Food Basket Delivery with the Stretch RE1

Jimin Sun, Prasoon Varshney
{jimins2, pvarshne}@cs.cmu.edu
Language Technologies Institute
Carnegie Mellon University

Abstract— **Injuries, frailty, and physical disabilities cause people to have trouble performing Activities of Daily Living (ADLs), including fetching food for themselves. A growing population of senior citizens in elderly care facilities and a shortage of staff pose logistical issues for frequent and timely deliveries of small objects customized to individual needs. Mobile manipulator platforms, such as the Hello Robot Stretch RE1, offer a solution for customized and autonomous robotic assistance. In this work, we present a system for point-to-point deliveries of food baskets using the Stretch RE1 Mobile Manipulator for patients in hospitals and nursing homes. We discuss the design decisions informed by technical challenges and stakeholder interactions and evaluate the feasibility and future work required for real-world deployments of sleek mobile manipulators through experimental evaluations. The code for our system is available on GitHub[1].**

I. INTRODUCTION

The US Department of Health and Human Services published a report[2] in January 2023 revealing that more than 1.4 million individuals reside in over 15,500 medicare and medicaid-certified nursing homes across the country. The COVID-19 pandemic has further exacerbated the staffing shortages in these nursing homes. According to a survey[3] conducted by the American Health Care Association and National Center for Assisted Living, 94% of nursing home providers faced staffing shortages during the summer of 2021, while 81% of assisted living communities experienced similar shortages. Certified nursing assistants, direct caregivers, and dietary staff were among the hardest-hit professions. The workforce situation has only worsened since 2020, with nearly 75% of nursing homes and almost 60% of assisted living communities reporting a decline in staff.

In the United States, mobility difficulties are the most prevalent disability among older adults (ages 65+), affecting over 15% of older adults (ages 65–74), 26% of those ages 75–85, and 48% of those aged 85 and above [1], [2], who often require care in assisted-living facilities. The Aging Concerns, Challenges, and Everyday Solution Strategies (ACCESS) [3] study has shown that individuals with prolonged mobility disabilities seek solutions that provide them autonomy with activities of daily living (ADLs) and without causing inconvenience to others.

To address these challenges, automated solutions in the form of service robots that can assist with ADLs are needed. Mobile manipulators, such as the Stretch Robot [4], offer a sleek design and a versatile array of sensors that enables audio-visual perception, allowing them to assist ADLs based on voice commands and complex tasks such as object detection, navigation, collision avoidance, and patient pose estimation. Moreover, the Stretch Robot provides a library of navigation, localization, and arm movement primitives through the open-source Robot Operating System (ROS) [5].

In this work, we focus on the application of point-to-point food basket deliveries, with a target population of senior citizens and those with disabilities living in nursing homes and elderly care facilities. We make the following contributions:

1) We describe a system design for utilizing the Stretch RE1 mobile manipulator and ROS to perform point-to-point food basket deliveries.
2) We outline the challenges and decision factors for real-world deployments in hospitals and nursing facilities based on interactions with stakeholders in the healthcare system.
3) Limitations and future work required to ensure the safety and effectiveness of such a system for use in nursing facilities are identified.

II. RELATED WORK

A. Assistive Robots for Elderly Care

In recent years, there has been a growing trend of using robots in healthcare to carry out tasks such as patient care, cleaning, and logistics [6]. With the healthcare sector facing a shortage of workers, service robots have emerged as a potential solution to help address this issue, while also improving patient outcomes, reducing costs, and increasing efficiency. In the field of healthcare robotics, there are various types of service robots available, including autonomous mobile robots, telepresence robots, and robotic exoskeletons. These robots have been employed for various purposes, such as sterilization, cleaning, COVID-19 testing, logistics, social care, and telehealth. For this project, we specifically focus on using service robots to assist with the logistics of delivering meals in healthcare facilities.

Various studies have utilized the Stretch RE robot [4] to explore the potential of assistive robots in helping humans with everyday activities, including dressing [7], [8], adjusting bedding [9], [10], and drinking [11]. A recent study claims the limitations of traditional design methods in the context of

[1] `https://github.com/prasoonvarshney/stretch-robot`
[2] US Department of Health and Human Services Report on Nursing Homes
[3] AHCA/NCAL Survey on Staffing Shortages Across Nursing Homes

Fig. 1. Industry Solutions: Left-to-Right: (1) Aethon TUG, (2) Moxi, (3) Relay, (4) Stretch RE1

Fig. 2. Candidate Basket Options

designing robotic interactions with older adults and suggests a collaborative design process with older adults in their own living environments [12].

B. Industry Solutions

Assistive robots have been productized and used in the industry to assist logistics in healthcare settings. One example is the Aethon TUG[4], which is designed to deliver medications and laundry in hospitals. While it can navigate different floors autonomously, its bulky size (22.4 x 33.8 x 47.7 inches) often makes it challenging to move through crowded corridors, requiring healthcare workers to manually restart it. Similar robots, such as Moxi[5] or Relay[6], are also relatively large or lack an end effector that enables flexible interaction with objects.

The Stretch RE1 offers a complementary solution with a leaner size (13.1 x 13.4 x 56.0 inches), making it easier to navigate through obstacles in crowded areas. This makes Stretch ideal for point-to-point custom deliveries of lightweight objects like water bottles, food items, and bags, whereas, the larger robots are better for bulk deliveries like laundry. In addition to its smaller size, the Stretch RE1 offers a built-in microphone and speaker, enabling flexible interaction with humans through voice commands. As opposed to the pre-programmed behavior and control through centralized fleet management software seen in products like the TUG, features of Stretch provide greater customization of robot operations.

Another dimension that makes Stretch an attractive option for widespread adoption in nursing homes is its relatively low cost of $19,950, compared to over $100,000 for various versions of Aethon TUG, and the expensive monthly subscription model for Moxi and Relay, which makes the latter robots affordable only to large hospital chains.

III. Methodology

In this section, we describe the process we followed to build our basket delivery system. We first look at the various design decisions informed by stakeholder challenges, technical challenges, and feasibility. Then, we dive into the details of the implemented navigation and object manipulation stacks. Lastly, discuss the limitations and ways to improve the current system.

A. Design Decisions

*1) **Stakeholder Interactions**:* We contacted two stakeholders working in the healthcare industry for feedback on the project direction. One is a nurse at the University of Pittsburgh Medical Center (UPMC), and the other is a doctor at BrightSpring Health Services.

Both stakeholders confirmed that a sleek mobile robot would be useful for navigation in small spaces. One stakeholder reported using Aethon's TUG for laundry and medication delivery in their workplace but noted the inconvenience of manual restarts when the robot gets stuck in cluttered corridors. Our proposed solution, based on the Stretch, would face these issues less frequently due to its smaller size. We also validated that the point-to-point delivery is adequate for hospitals as it accommodates the different meal needs of patients. Both stakeholders also agreed that a voice control interface would be useful. Additionally, one stakeholder commented that the robot should ideally be able to adapt to changes in the patient's position upon delivery.

Overall, based on the feedback received from these stakeholders, we believe that a compact, mobile robot with a voice control interface and the ability to adapt to changes in patient position would be an ideal solution for hospital logistics. While we focus on the delivery component in this project, we hope to see these aspects integrated into future robot designs to enhance their usefulness in healthcare settings.

*2) **Choice of Basket**:* Figure 2 illustrates the different basket options we considered for our project, where the leftmost basket was chosen as the final basket for the robot. Initially, we considered a flatter food tray, commonly utilized in cafeterias. However, we found that the flat container was incompatible with the existing end effector, as Stretch features only a single end effector. Attempting to hold a plate-like object with a single hand resulted in an unstable position, causing the content to shift even with meticulous human teleoperation.

As a substitute, we opted for a basket-shaped container with a handle on the top that can be easily grasped by Stretch's end effector. This shape provides stability to the container during pick up and transportation, even when using one hand, as the handle is located at the center of mass. In addition, the depth of the first basket we utilized provides

[4]https://aethon.com/brochures/
[5]https://www.diligentrobots.com/moxi
[6]https://www.relayrobotics.com/

Fig. 3. A map of the JPMorgan Chase AI Maker Space at the Tepper building at Carnegie Mellon University

a degree of privacy for the items it contains. Furthermore, the basket's fabric material ensures a safer interaction with humans compared to the second plastic container, which has relatively sharper edges.

*3) **ArUco Tag Detection**:* For an intelligent system to be effective, it must be capable of detecting the object of interest and moving its joints to accessible locations without relying on pre-programmed coordinates. In this study, we aimed to improve the localization and grasping components of our robot's basket delivery system by comparing two detection methods – object detection and ArUco marker detection.

As the Stretch ROS system comes equipped with a learned perception stack, we initially tested the YOLO v3 object detection network on our basket. However, as our basket is not a typical prototype of a basket, the detection model struggled to identify it correctly, often misclassifying it as other objects, such as "cake."

As a result, we opted to leverage ArUco markers by attaching the markers to the basket. While fine-tuning the object detection network would have been an alternative option, we chose ArUco markers as they allow for rapid detection from a particular viewpoint and offer superior robustness in terms of location estimation when compared to object detection models.

*4) **Assumptions**:* The assumptions made throughout the project are as follows:

1) The basket possesses a handle on its top to enable easy grasping.
2) The map of the room is mapped out in advance, for the robot's navigation stack to load.
3) A human is available to help the robot with its initial localization.
4) The coordinates of the patient bed are known and fixed.
5) While the coordinates of the basket remain unknown, it is likely that the basket resides in the kitchen area.
6) An Aruco tag is affixed to the basket.

B. Navigation Stack

*1) **Mapping**:* We use the ROS navigation stack and the Lidar sensor to map out the entire room of the robot's operation in advance. Figure 3 shows an image of a built map of the AI Maker Space at Carnegie Mellon.

*2) **ArUco Tag Detection**:* For detecting ArUco tags, we utilized the following ROS launch files, which are provided as part of the `stretch_core` API: (1) `stretch_driver`, (2) `d435i_low_resolution`, (3) `stretch_aruco`. These files enable the *detect_aruco_markers* node, which publishes translation and rotation coordinates of the detected tag with respect to the map and odometry frames onto the `tf` ROS topic.

We generate a 6×6 ArUco marker with ID 0 using an online tool.[7] Figure 4 shows an example of an ArUco tag. The marker was then printed in a 10×10 (cm) size and tagged onto the basket. It is worth noting that the dimensions of the marker before and after printing may differ. Therefore, we carefully measured the printed ArUco tags and entered their lengths into the `stretch_marker_dict.yaml` configuration file. This step is crucial in enabling the robot to estimate the depth of the basket accurately based on the known absolute size and the relative size observed in the RGB sensory input.

The following two steps are repeated until the target ArUco tag is found:

- Lookaround: The realsense camera is rotated 360 degrees in 45-degree increments.
- Known Priors: The baskets are expected to be near the kitchen area depicted in Figure 3. A navigation goal to predefined coordinates is sent to continue the search process.

The searching loop ends as soon as an ArUco tag is detected, or if the max retries count is reached.

[7] `https://chev.me/arucogen/`

*3) **Reaching to ArUco:*** Once the ArUco tag has been detected, the ROS navigation stack is used to query the coordinates (x, y, z) and the orientation quaternion q of the target object relative to the *map* frame.

While the queried coordinates and orientations are in the 3D space, they are first projected down to the $z = 0$ plane for computing point navigation goals, since the robot base is on the ground. The projection follows involves computing euler angles along X, Y, and Z axes from the orientation quaternion, and retaining only the angle along the Z axis, θ.

1) The euler angle ϕ along the Z-axis for the surface normal projecting out of the ArUco tag is given by:

$$\phi = \theta + 3 * \pi/2$$

and, the the euler angle γ along the Z-axis for the desired base_link direction is given by:

$$\gamma = \phi - \pi/2 = \theta + \pi$$

Note that the angles ϕ and γ are shifted into the range $[-\pi, \pi]$ using a simple modulo operation.

2) The $X-Y$ offset in the direction of the surface normal projecting out of the ArUco tag is then given by the tuple:

$$(R * cos(\phi), R * sin(\phi))$$

The wrist alignment offset opposite to desired base direction, to address the relative gap r between the center of the Dex Wrist and robot arm lift is given by:

$$(-r * cos(\gamma), -r * sin(\gamma))$$

3) The final $X-Y$ 2D goal coordinates to specify the PointNav goal for the robot are then computed as:

$$x' = x + R * cos(\phi) - r * cos(\gamma)$$

$$y' = y + R * sin(\phi) - r * sin(\gamma)$$

where R is kept constant at 40cm informed by the `wrist_extension` for manipulation actions, and r is measured to be 8cm.

4) The arm joint lift l is informed by the height of the basket based on the queried z co-ordinate of the basket and an offset c = 5cm measured as the height difference between the ArUco tag and the handle of the basket.

$$l = z + c$$

C. Manipulation Stack

The manipulation actions used are based on a sequence of hard-coded joint trajectory goals for the arm, wrist, and gripper. Once the Stretch robot reaches nearby the basket in the correct orientation, it performs the `PICK` action, navigates to the hard-coded coordinates for the patient bed, and performs the `PLACE` action.

In our current implementation, the `PICK` action executes the following sequence of low-level actions:

1) Lift the arm to the right height (based on the computed joint lift l in Section III-B.3)

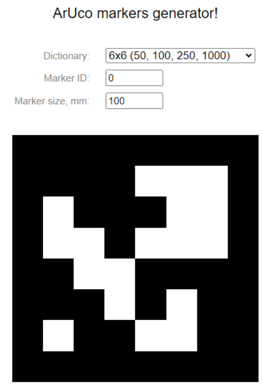

Fig. 4. Sample ArUco Marker generated using an online tool

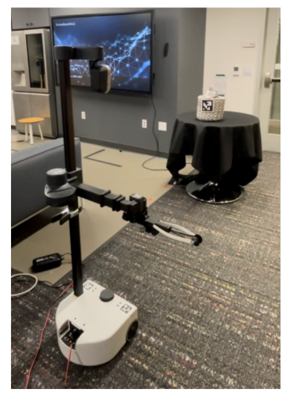

Fig. 5. Stretch robot in its initial position for the 8 end-to-end experiments

2) Move the wrist slightly up (to decrease chances of collisions while extending the arm) and open the gripper
3) Extend the arm based on the constant R in section III-B.3
4) Move the wrist down across the basket handle and close the gripper
5) Retract the arm, lift it to maximum height, and fold the wrist

The `PLACE` action performs the following sequence of low-level actions:

1) Extend the wrist and lower the arm based on hardcoded patient bed or couch measurements
2) Move the wrist down until contact and open the gripper
3) Lift the arm to max height
4) Retract the arm and fold the wrist

D. Limitations and Future Work

Navigation: Our approach utilizes a static map of the environment, which means that the robot cannot avoid obstacles autonomously or adaptively re-plan its route based on the most recent sensory information. This would pose a significant challenge when deploying the robot in real-world settings where dynamic objects (e.g., carts) or humans are present, and there is a need for robots to navigate around them.

Step	Success Rate (%)
Basket Detected	100
Navigation to Basket	87.5
Pick up Basket	87.5
Navigation to Patient Bed	75
Place Basket	62.5

TABLE I
OVERALL SUCCESS RATE FOR END-TO-END BASKET DELIVERY ACROSS EIGHT TRIALS.

Step	Duration (sec)
Nav to Basket	20.99
Navigation to Basket	44.02
Pick up Basket	51.71
Navigation to Patient Bed	35.11
Place Basket	53.32
Total	205.15

TABLE II
AVERAGE TIME TAKEN AT EACH STEP FOR END-TO-END BASKET DELIVERY ACROSS EIGHT TRIALS.

Manipulation: Our pre-programmed `PICK` and `PLACE` actions are effective for baskets with a standardized form-factor. However, failure cases may occasionally occur due to the margin of error in orientation and position of the navigation goals, which can deviate by a few degrees and centimeters, respectively. The mean and standard deviations of these errors are presented in Table III based on the results of our experiments. Currently, the robot is unable to detect when it fails to grasp an object. Consequently, it proceeds with subsequent plans, navigating to the goal location even when its gripper is empty. Incorporating a method that enables the robot to recognize its failures through sensory information, such as a camera image of its gripper or the torque sent to the gripper/arm's motor, and retry based on that information would create a more autonomous agent that humans could reliably use.

Localization: Another limitation lies in the initial pose estimation step. The Stretch robot has two options for determining its pose: an explicit estimate at the start of each run or self-estimation by moving around. However, we found that the latter option was not viable due to battery issues with our specific robot. In our experiments, we provided initial pose estimates manually in Rviz Graphical User Interface at the start of each run or started the robot at the same origin location in the room for each run. However, both options hinder easy deployment in real-world settings. We believe that future iterations of the Stretch robot require improved hardware sensors and software capabilities for localization to overcome this limitation.

IV. EXPERIMENTAL EVALUATIONS

A. End-to-End Delivery with Aruco Tag Detection

We performed eight end-to-end trials with the current system design. Figure 5 shows the Stretch Robot we used in its initial configuration after performing manual pose estimation.

1) Success Rates: Table II presents the end-to-end success rates of the high-level tasks executed by our system. A failure in any of the steps implies the subsequent steps are also considered unsuccessful. While the basket is detected successfully at all trials, errors in pose estimation often cause the navigation and manipulation to fail. It is also noteworthy that the placing task, despite being hardcoded, failed once due to the basket handle getting entangled in the open gripper, causing the robot to pick the object up again. This issue can be addressed by implementing a weight sensor on the Dex Wrist and incorporating recovery behaviors.

Step	Translation Error (m)	Orientation Error (deg)
Navigation to Basket	0.07 (std=0.01)	3.5 (std=1.01)
Navigation to Patient Bed	0.06 (std=0.02)	1.62 (std=1.34)

TABLE III
OVERALL SUCCESS RATE FOR END-TO-END BASKET DELIVERY ACROSS EIGHT TRIALS.

2) Duration: Furthermore, we conducted a time analysis to measure the duration of each step involved in completing the task. Our results show that the most time-consuming steps for the robot are picking up the basket and placing the basket in the patient's bed area. This is due to the current approach of sending joint trajectory goals sequentially, with each joint being moved one at a time (e.g., extending the arm, lowering the arm, opening the gripper). This design was chosen to enhance safety during the transfer of objects to the patient. Moving one joint at a time provides a more predictable movement from the perspective of the human user, as rapidly extending and lowering the arm could potentially result in a collision with the patient. We believe that this trade-off between speed and safety can be improved in future iterations of the robot.

3) Failure Cases: Outside of the eight trials, we observed several instances of failures in picking up the basket that

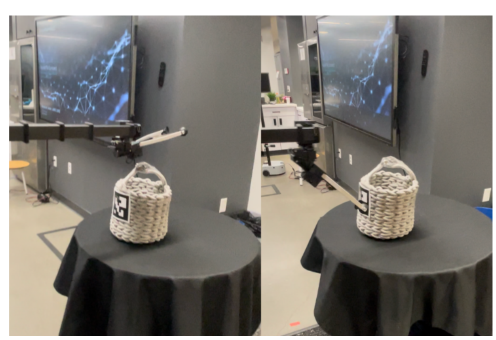

Fig. 6. Arm incorrectly extended to the right of the basket as the wrist drops down in an attempt to pick it up

led us to adjust the distance r between the end-effector and the wrist, as mentioned in Section III-B.3. An example of this error is illustrated in Figure 6, where the wrist deviated slightly downwards to the right side of the basket.

4) Navigation Errors: Throughout the eight end-to-end trials, navigation demonstrated an average translation error of 6-7 centimeters and an average orientation error of 2-4 degrees, as presented in Table III. Notably, despite these marginal errors, the pick action achieved a 100% success rate when the robot was able to navigate to the basket successfully. An interesting observation was that the robot was able to hold the basket from the side instead of the handle during one trial. The basket's sturdy but deformable structure played a role in contributing to the overall success rate of the `PICK` action.

V. Conclusion

This study presents a system design utilizing a Stretch Robot for point-to-point deliveries of small daily objects carried by a basket. The Stretch robot possesses a sleek form-factor and the advantage of navigating through cluttered hallways and small spaces, whereas the size of currently deployed robots often restricts their ability to navigate effectively. The Stretch robot's arm enables point-to-point deliveries, and our experiments achieved a composite success rate of 62.5% across eight trials. While the current system's overall success rate for end-to-end deliveries might not be sufficient for deployment, we identify key potential improvements, including utilizing sensors on the end-effector for more precise grasping and placing of objects.

References

[1] A. W. Roberts, S. U. Ogunwole, L. Blakeslee, M. A. Rabe, *et al.*, *The population 65 years and older in the United States: 2016*. US Department of Commerce, Economics and Statistics Administration, US ..., 2018.

[8] A. Kapusta, Z. M. Erickson, H. M. Clever, W. Yu, C. K. Liu, G. Turk, and C. C. Kemp, "Personalized collaborative plans for robot-assisted dressing via optimization and simulation," *Autonomous Robots*, vol. 43, pp. 2183 – 2207, 2019.

[2] W. Erickson, C. Lee, and S. Von Schrader, "Disability statistics from the american community survey (acs)," *Ithaca, NY: Cornell University Yang-Tan Institute (YTI). Retrieved from Cornell University Disability Statistics website: www. disabilitystatistics. org*, 2017.

[3] L. M. Koon, E. T. Remillard, T. L. Mitzner, and W. A. Rogers, "Aging concerns, challenges, and everyday solution strategies (access) for adults aging with a long-term mobility disability," *Disability and Health Journal*, vol. 13, no. 4, p. 100936, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1936657420300613

[4] 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)*, 2022, pp. 3150–3157.

[5] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng, *et al.*, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.

[6] J. Holland, L. Kingston, C. McCarthy, E. Armstrong, P. O'Dwyer, F. Merz, and M. McConnell, "Service robots in the healthcare sector," *Robotics*, vol. 10, no. 1, 2021. [Online]. Available: https://www.mdpi.com/2218-6581/10/1/47

[7] Z. Erickson, H. M. Clever, G. Turk, C. K. Liu, and C. C. Kemp, "Deep haptic model predictive control for robot-assisted dressing," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 4437–4444.

[9] K. Puthuveetil, C. C. Kemp, and Z. Erickson, "Bodies uncovered: Learning to manipulate real blankets around people via physics simulations," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1984–1991, 2022.

[10] K. Puthuveetil, S. Wald, A. Pusalkar, P. Karnati, and Z. Erickson, "Robust body exposure (robe): A graph-based dynamics modeling approach to manipulating blankets over people," 2023.

[11] H. Reddy, C. Jaroenkunathum, and A. Chen, "Aquabot: Assistive drinking robot for the mobility-impaired."

[12] L. Stegner, E. Senft, and B. Mutlu, "Situated participatory design: A method for in situ design of robotic interaction with older adults," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–15.