Proof-of-Concept: A Hands-Free Interface for Robot-Assisted Self-Feeding*

John R. Schultz,[1] Andrew B. Slifkin,[2] Hongkai Yu[3], and Eric M. Schearer[1]

Abstract— **Eating and drinking is an essential part of everyday life. And yet, there are many people in the world today who rely on others to feed them. In this work, we present a prototype robot-assisted self-feeding system for individuals with movement disorders. The system is capable of perceiving, localizing, grasping, and delivering non-compliant food items to an individual. We trained an object recognition network to detect specific food items, and we compute the grasp pose for each item. Human input is obtained through an interface consisting of an eye-tracker and a display screen. The human selects options on the monitor with their eye and head movements and triggers responses with mouth movements. We performed a pilot study with four able-bodied participants and one participant with a spinal cord injury (SCI) to evaluate the performance of our prototype system. Participants selected food items with their eye movements, which were then delivered by the robot. We observed an average overall feeding success rate of 89.1% and an average overall task time of 31.4±2.4 seconds per food item. The SCI participant gave scores of 90.0 and 8.3 on the System Usability Scale and NASA Task Load Index, respectively. We also conducted a custom, post-study interview to gather participant feedback to drive future design decisions. The quantitative results and qualitative user feedback demonstrate the feasibility of robot-assisted self-feeding and justify continued research into mealtime-related assistive devices.**

I. INTRODUCTION

According to the World Health Organization, there are more than 1 billion people, or 15% of the world's population, living with some sort of a motor impairment [1]. Motor impairments have been shown to have a significant effect on an individual's ability to perform activities of daily living (ADLs), such as bathing, grooming, and feeding [2]. Commercially available assistive robot arms can help people perform a wide range of ADLs on their own [3]. One of the most highly rated ADLs by people with movement disorders is the ability to prepare a meal and feed oneself [4]. Stand-alone devices that are specifically designed for robot-assisted feeding (RAF) include the Meal Buddy and the Obi Feeder [5], [6].

While these commercial solutions have proven effective in certain situations, they have a number of drawbacks. Assistive robotic arms are effective for general purpose reaching and grasping tasks, but can be difficult to control. RAF devices address this issue by constraining the application to a mealtime setting. This allows for the design of more user-friendly interfaces, but limits the usability of these devices in different environments. Additionally, people may need help to feed themselves due to a variety of movement disorders, requiring the interface to be either heavily customized for each individual, or robust enough to be usable by people with different ability levels. In this work, we address these needs by presenting methods for recognition, localization, and grasping of food items on a table, a bi-directional communication human-machine interface, and an initial evaluation of our prototype system. The quantitative results and qualitative user feedback will inform future design decisions to tailor the system to the needs of individuals with SCI.

II. RELATED WORK

Although there are many people who require help to feed themselves, the body of literature on robot-assisted feeding is relatively small. Perhaps the most notable RAF device is the Assistive Dexterous Arm from the Personal Robotics Lab [7]. This system consists of a wheelchair-mounted robot arm commonly used by people with movement disorders. Much of their research was focused on classification of food items and grasping strategies, particularly motion planning and utensil skewering forces [8].

Another group that developed a notable RAF system is the HealthCare Robotics Lab, approaching the problem from a caregiver's perspective [9]. Instead of a wheelchair-mounted robot arm, their design centered around a two-armed, mobile household robot. The system was capable of scooping deformable foods out of bowls, equipped with sophisticated perception modules, and controlled via a custom graphical user interface (GUI).

Both groups have investigated critical facets of the complex area of robot-assisted feeding. Our design centers around restoring function to the individual so that they are in control over their meal, rather than a caregiver or robot. The interface allows the user to communicate their intentions to the system with naturally available command inputs, while simultaneously receiving visual feedback from the device. Above all, the user is in control of the most important mealtime decisions, what to eat and when to eat it. By gathering feedback from individuals with SCI, we hope to uncover specific feeding tasks with which to drive the future development of our system, so that it may best address the needs of people with movement disorders.

*This work was supported by the Ohio Department of Higher Education

[1]John Schultz and Eric Schearer are with the Center for Human Machine Systems at Cleveland State University, Cleveland, OH 44115, j.r.schultz12@csuohio.edu, e.schearer@csuohio.edu

[2]Andrew Slifkin is with the Department of Psychology at Cleveland State University, Cleveland, OH 44115, a.slifkin@csuohio.edu

[3]Hongkai Yu is with the Department of Electrical Engineering and Computer Science at Cleveland State University, Cleveland, OH 44115, h.yu19@csuohio.edu

III. METHODS

A. System Overview

The goal of our RAF system is to allow individuals with movement disorders the ability to feed themselves. With this in mind, the system consists of the following primary components:

1) A method for perceiving and localizing food items
2) A mechanism to obtain and deliver the food items to the human
3) An interface which facilitates bi-directional communication between the human and the device

The robot used in this work was the Baxter Research Robot (Rethink Robotics; Bochum, Germany). We fixed an L515 LIDAR camera (Intel; Santa Clara, California) on the robot's wrist using a custom, 3D-printed mount. A tablet monitor was mounted to the table surface fit with a Tobii Eye Tracker 4 device (Tobii; Danderyd, Sweden). We developed a custom GUI application that displays the output from the depth camera as well as different control options. The user moves their eyes to select options on the tablet which controls the robot to pick up and deliver the selected food items. Internal communication was handled by the Robot Operating System (ROS). An overview of the RAF system is shown in Figure 1.

B. Object Recognition

To recognize food items on a plate, we used a segmentation mask, regions-based convolutional neural network (maskR-CNN) based on a ResNet+FPN backbone from detectron2 [10]. The network was trained to detect four object classes: 1) Carrot, 2) Celery, 3) Pretzel, and 4) Robot Gripper. These items were chosen due to their uniform shape and non-compliant material to reduce the complexity of grasping for this proof-of-concept. The network takes the depth camera's RGB data stream as an input and returns bounding boxes and instance level segmentation masks of each detected object (Fig. 2). We collected and annotated 987 images for the training set and 267 images for the test set. Details regarding the dataset and training parameters can be found in Table I. The network was trained on a GeForce RTX 3090 GPU (NVIDIA; Santa Clara, California). Average Precision (AP) values closer to 100.0 indicate better detection performance.

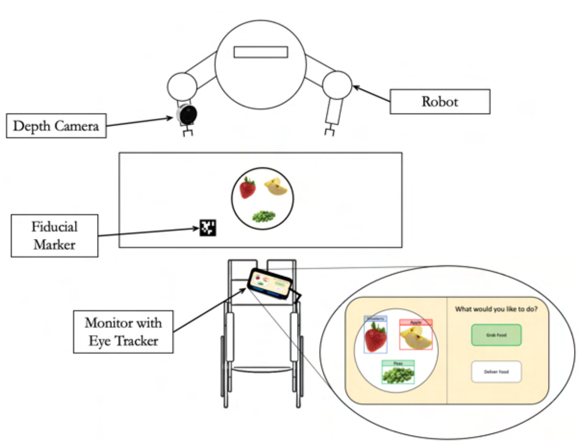

Fig. 1. Top-Down view of the robot-assisted feeding system. The human is seated across the table from the robot. The user moves their eyes and head to select options on the monitor which control the robot to interact with food items.

TABLE I
OBJECT DETECTION DATASET DETAILS

Class	**Annotations**	**AP**
Carrot	1714	86.2
Celery	1383	92.0
Pretzel	1785	89.6
Gripper	598	86.6
Training Parameters	Max Iterations: 50,000 Batch Size: 512 Training Time: $\approx$ 10 hours	

C. Table Plane Registration

We placed an AprilTag fiducial marker on the surface of the table to define the table's coordinate frame (Fig. 2). The tag is detected by the AprilTag 3 ROS package [11]. We then performed a Direct Linear Transformation (DLT) to map camera pixels $(u,\ v)$ to table coordinates in meters $(x,\ y)$ [12]. In this method, the corners of the tag are used as correspondence points. If the location of n correspondence points are known in both the camera frame and the table frame, the solution to a set of linear equations is $2n$ DLT parameters P (1). These DLT parameters are then used to transform any camera pixel location to planar table coordinates (2), (3).

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 & u_1x_1 & u_1y_1 \\ 0 & 0 & 0 & x_1 & y_1 & 1 & v_1x_1 & v_1y_1 \\ x_2 & y_2 & 1 & 0 & 0 & 0 & u_2x_2 & u_2y_2 \\ 0 & 0 & 0 & x_2 & y_2 & 1 & v_2x_2 & v_2y_2 \\ x_3 & y_3 & 1 & 0 & 0 & 0 & u_3x_3 & u_3y_3 \\ 0 & 0 & 0 & x_3 & y_3 & 1 & v_3x_3 & v_3y_3 \\ x_4 & y_4 & 1 & 0 & 0 & 0 & u_4x_4 & u_4y_4 \\ 0 & 0 & 0 & x_4 & y_4 & 1 & v_4x_4 & v_4y_4 \end{bmatrix} \cdot \begin{bmatrix} P_1 \\ P_2 \\ P_3 \\ P_4 \\ P_5 \\ P_6 \\ P_7 \\ P_8 \end{bmatrix} = - \begin{bmatrix} u_1 \\ v_1 \\ u_2 \\ v_2 \\ u_3 \\ v_3 \\ u_4 \\ v_4 \end{bmatrix} \tag{1}$$

$$x = \frac{(P_2 + uP_8)(P_6 + v) - (P_5 + vP_8)(P_3 + u)}{(P_1 + uP_7)(P_5 + vP_8) - (P_4 + vP_7)(P_2 + uP_8)} \tag{2}$$

$$y = \frac{(P_4 + vP_7)(P_3 + u) - (P_1 + uP_7)(P_6 + v)}{(P_1 + uP_7)(P_5 + vP_8) - (P_4 + vP_7)(P_2 + uP_8)} \tag{3}$$

D. Camera to Robot Calibration

The custom, 3D-printed mount which fixes the depth camera to the robot's wrist allows for tuning of the camera's angle (Fig. 3). To combat the need to frequently re-define the camera pose, we developed a calibration procedure. First, a food item such as a pretzel is placed on the table within view of the camera. The centroid of the food item in table coordinates $(x,\ y)$, as well as the depth value (z) returned

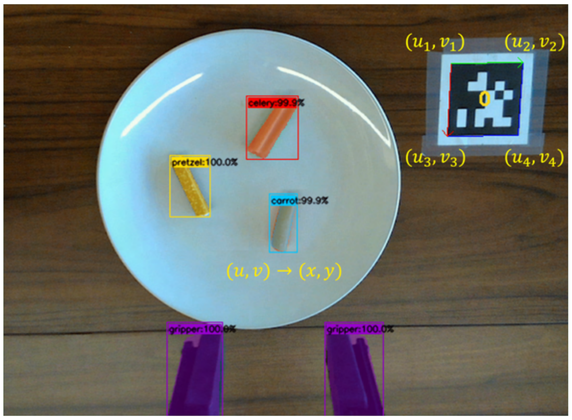

Fig. 2. The object detection network returns bounding boxes and instance segmentation masks for each food item. The fiducial marker is used to transform the pixel coordinates into planar table coordinates.

from the depth camera is saved as O_1. Then, the robot's gripper is manually moved to the center of the food item and the position of the gripper in the robot's frame (X, Y, Z) is saved as O_2. Finally, the transformation matrix (T) between the camera's frame and the robot's frame is computed using the methods described in Söderkvist, et al. for a minimum of 3 points [13]. Once the calibration has been completed, any (x, y) coordinate pair in the table coordinate frame can be transformed into robot coordinates (X, Y, Z) by multiplying by the transformation matrix (T).

E. Grasping Food Items

To properly grasp the food item with the robot's two-finger gripper, we define a grasp pose. In a recent study, Gallenberger et al. found that grasping long, slender food items toward one end rather than in the center made it easier for people to bite the food item [14]. Therefore, we defined the grasp point as halfway between the centroid and the bottom edge of the food item. To find the grasp point, we first compute the rectangle with minimum area that circumscribes the food item's segmentation mask using the OpenCV function `minAreaRect()`. Once the rectangle has been obtained, it is trivial to compute the grasp point. The angle that the long axis of the rectangle makes with the u-axis of the camera frame defines the rotation angle. The grasp point and the rotation angle are transformed to the robot's frame using the methods described in section III-D (Fig. 4).

F. System Operation and Human-Machine Interface

1) Overview: We used a 15.6" monitor to display a custom GUI developed using Qt Creator and QML (The QT Company; Espoo, Finland). The GUI primarily displays the output from the depth camera's RGB data stream overlaid with custom graphics produced by OpenCV (Fig. 5). When food items are detected, the GUI draws bounding boxes around each food item. The user is then responsible for selecting a food item with their eye movements. Once selected, the robot acquires the food item. The user opens their mouth to trigger the food item transfer. The robot then slowly approaches the user's mouth and releases the food item. This process is repeated for all the food items on the plate.

2) Food Item Selection: Mounted on the bottom of the monitor is a Tobii Eye tracker. We used Talon, a hands-free input replacement program to interface with the eye tracker [15]. Talon's `controlMouse()` program allows the user to control a mouse cursor with their eye and head movements. Instead of a traditional mouse cursor, we displayed a custom, interactive cursor. When the cursor remains stationary over a selectable item for a duration, or dwell time, the cursor's interior changes color in a rotational sweeping motion similar to a radar scanner. When the animation is finished and the cursor is "filled up", the target food item is selected. In this way, the user receives feedback for when the system makes a selection and allows time for the user to stop the selection process if they wish.

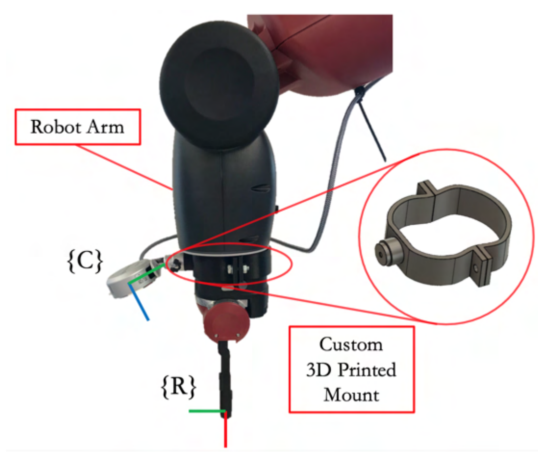

Fig. 3. The depth camera is mounted to the robot's wrist. A calibration is performed to describe the camera coordinates C in the robot's coordinate frame R.

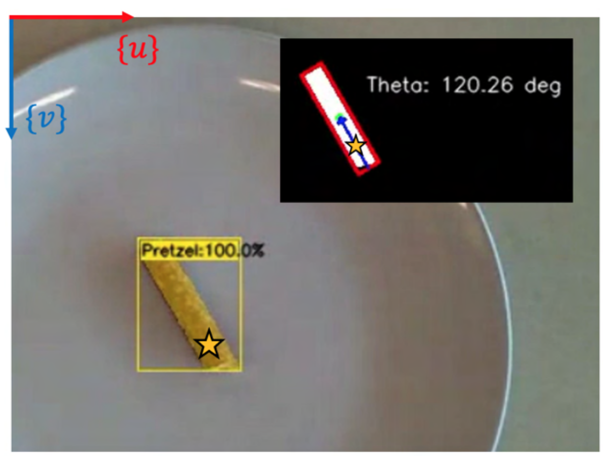

Fig. 4. The grasp position (yellow star) is defined as the 3D location of the point halfway between the centroid and the bottom center of the food item. The grasp angle (blue arrow) is defined as the angle the long side of the rectangle makes with the u-axis of the camera frame.

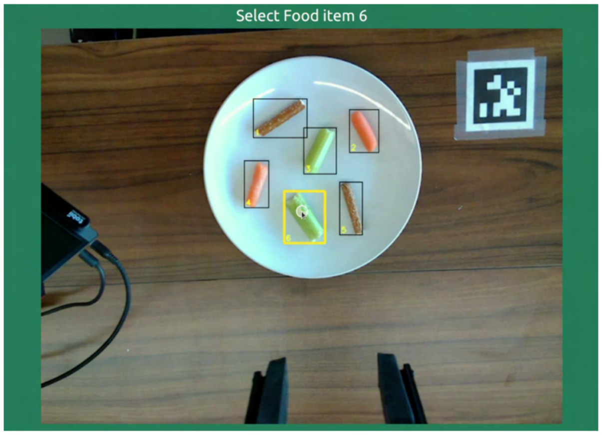

Fig. 5. Screenshot of the Item Selection screen on the GUI. Each food item is assigned a number. The participant moves their eyes to select a food item chosen at random. After a dwell time, the item is selected.

The cursor only animates when it is stationary within the bounds of a selectable item. For example, when the user directs their gaze to a food item, the bounding box is highlighted, indicating the item is selectable. If the cursor remains stationary in the bounding box for a dwell time (1 second), the cursor animation triggers, and the food item is selected. We found that this combination of user input and visual feedback was an appropriate trade-off between autonomy and control (Fig. 5).

3) Food Item Acquisition: Once the food item has been selected, the robot computes the grasp pose using the methods described in section III-E. As the robot control was not the primary focus of this proof-of-concept, we simply used Bazter's built-in Cartesian endpoint controllers to acquire the food items with the computed grasp pose.

4) Food Item Transfer and Facial Keypoint Detection: To detect the user's facial keypoints, we used the `face-alignment` python package [16]. This package detects 60 facial landmarks, 12 of which belong to the mouth. We fit an ellipse to these 12 landmarks to detect when the mouth was open versus closed, set to a comfortable value for each participant (Fig. 6). When the user opens their mouth, the robot approaches and releases the food item after a delay of 1.5 seconds [14].

G. Evaluation

We had four able-bodied individuals (2 Male, 2 Female; 24-26 years old) and one individual (Male, 40 years old) with a C4 level spinal cord injury evaluate the system. We recorded success rates and task times for each participant. The SCI participant completed the System Usability Scale (SUS), the NASA Task Load Index (TLX), and a custom, post-study survey to qualitatively asses the system [17], [18]. For safety, the food item transfer position was defined before data collection to ensure no collisions with the participant. Experiment protocol was reviewed and approved by the Cleveland State Institutional Review Board (IRB #FY2021-272).

Fig. 6. Visualization of mouth detection. Facial landmarks and ellipse are shown in red for when the mouth is closed and green when the mouth is open. Opening the mouth triggers the robot to begin the item transfer sequence.

For the duration of the study, the participant was seated in a powered wheelchair at a table with the robot across from them (Fig. 7). A plate of 6 food items was placed on the table, two of each carrot, celery, and pretzel. The food items were randomly arranged on the plate and sparsely distributed to simplify grasping. The monitor was mounted on the table to the participant's right. The participant performed a 30-second eye tracker calibration and was allowed 5 minutes to get acclimated to controlling the cursor with their eyes and selecting options on the GUI. The monitor displayed bounding boxes around the food items, which were numbered from 1 to 6. A message was displayed on the monitor prompting the participant to select a food item in a random order (Fig. 5). After the food item was selected, the robot reached for and grasped the food item. Then, the participant was prompted to open their mouth when they were ready to accept the food item. Once triggered, the robot slowly approached the participant's mouth and released the food item. This process was repeated for each food item on the plate. The process of selecting, acquiring, and delivering a food item was considered one trial.

Each participant performed approximately 60 trials. If the participant was able to accept the food item without dropping it, the trial was considered successful. If not, the trial was unsuccessful and the reason for the failure was recorded.

IV. RESULTS

A. Success Rates and Task Times

On average, we observed an overall success rate of 89.1%, meaning the participant was able to eat the food item the robot acquired. Out of the 270 total attempted trials, 22 carrots, 5 celery, and 2 pretzels were unsuccessfully acquired. All of the able-bodied participants displayed a selection success rate of 100%, meaning, they used their eye movements to select the indicated food item in every instance. Interestingly, the SCI participant selected the correct item only 75% of the time. In the post-study survey, he mentioned that he did not pay much attention to the instructions on the monitor, rather, focusing more of his attention on the robot as it was moving. Therefore, we do not believe that the lower selection success rate was due to an inability to operate the eye tracking system, but rather a priority of attention focus.

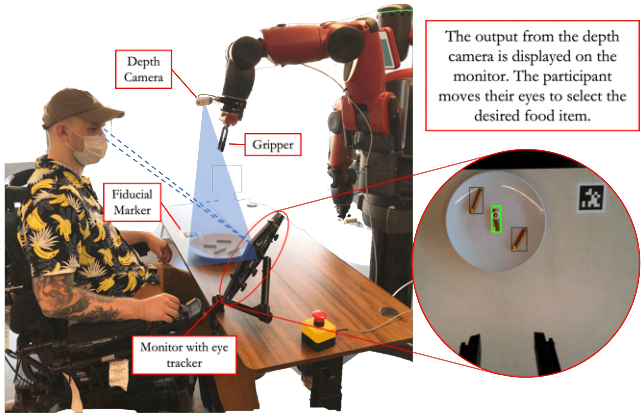

Fig. 7. Experimental Setup. The participant selects the desired food item with their eye movements. The robot then acquires, delivers, and transfers the food item to the participant.

Although this situation is not relevant in application, it is an interesting observation and may require further investigation in future studies.

The participants were able to open their mouths to trigger the release of the food item in every instance. The average trial completion time, from when the participant was prompted to select an item, to when the food item was released, was 31.4±2.4 seconds. Of this, 25.9 seconds, or 82% of the total time, was occupied by robot movement. This means that, on average, user input was required for 5.5 seconds per food item.

B. Post-Study Evaluation

In addition to measuring success rates and task times, we gathered feedback from the SCI participant on his experience with the system directly following data collection. The SUS assesses the participant's perceived usefulness of the system, while the TLX assesses the participant's self-reported level of exertion required to complete the task. The purpose of the custom post-study survey was to gather open-ended, subjective feedback and criticism from the participant about their experience with the system. The participant gave scores of 90/100 and 8.3/100 on the SUS and TLX, respectively. On the SUS, a higher number is preferred and a score above 80 is considered excellent. On the TLX, a lower number is preferred and represents minimal effort required to complete the task. These results and the outcomes of the post-study survey are discussed further in section V-B.

V. Discussion

A. System Performance

Ideally, the success rate of our RAF system would be 100%. However, all of the failed trials resulted from improper grasping of the food items, with the majority being carrots. We believe this is due in part to a lack of effective friction between the food item and the gripper as well as slight errors in the localization of the food items. The rounded, smooth, and sometimes slippery surface of the carrots and celery made it difficult for the robot's gripper to maintain an effective grasp. This issue was exacerbated if the depth estimation was imprecise. For this reason, we believe improper grasping to be the primary source of success rate error in this study.

While they do not report success rate for an entire feeding trial, the Assistive Dextrous Arm from the Personal Robotics Lab demonstrated an acquisition success rate of 70% [14]. Because all of our failure points occurred during acquisition or transfer, this is a relatively appropriate comparison. That being said, the experiment that Gallenberger et al. presented included a greater number of food item classes, as well as the more complex task of skewering food items with a fork. Our study was intentionally designed to limit the complexity of the grasping task to devote more emphasis on the communication interface and overall feeding task.

The average task time per food item was 31.4 seconds, with 82% of this time occupied by robot movement, and 5.5 seconds required for user input. The literature does not present an evaluation of task times, however the SCI participant mentioned that the user input mechanism and amount of time required was acceptable. He mentioned that it would be preferable to shorten the overall task time, or better yet, to have control over the movement speed of the robot. This could be addressed in future studies by including a different form-factor robot more appropriately designed for a feeding task, as well as trajectory optimization for movement duration.

B. SCI Participant Feedback

The SCI participant scored the RAF system highly on the SUS. He mentioned that if it was able to help him with even one feeding task, he would be willing use it. He did note that if food items were consistently dropped, it may be enough for him to abandon the system. However, because the food items included in this study were unlikely to cause a mess compared to food like sandwiches, soup, etc., he said the failures were less of a concern.

In addition to the SUS, the SCI participant reported a minimal effort required to operate the RAF system, as indicated by his score on the TLX. He mentioned the most difficult aspect for him was neck strain due to the placement of the tablet monitor. This could be mitigated by a more appropriate mounting mechanism customized to the individual.

In the post-study interview, the participant mentioned that he felt the system had the potential to address his feeding needs. He mentioned that the eye-tracking and mouth trigger input mechanisms were appropriate, but suggested giving the user control over the robot's speed. Other tasks he envisioned the system performing were holding poker cards, brushing his teeth, and cutting his hair. Finally, he remarked that caregivers rely heavily on non-verbal cues such as eye movement and head nods to know when and how to feed him. He also mentioned that caregivers tend to take food away from him before he has a chance to complete a bite. Future research should investigate the complex relationship between the caregiver and the individual during food item hand off.

C. Limitations

One apparent limitation of the current proof-of-concept demonstration was the intentional choice of easy-to-grasp food items. The primary goal of this work was to understand the challenges of feeding from the perspective of an individual with an SCI, and how to best develop an RAF system to meet these needs. By gathering feedback from participants early in the development process, we hope to drive future design decisions toward addressing the real needs of people who have difficulty feeding themselves.

Another limitation of this work was the small sample size. Having able-bodied individuals test the system was purely for task repetitions to gather a better assessment of the system's grasping performance. We believe the real value comes from having members of the beneficiary community test the system and provide feedback on its usefulness. Future studies should include large-scale qualitative assessments outlining the current feeding challenges of the SCI community and how best to address them with assistive feeding technologies.

Finally, the task scope of the current study was limited. We included only four classes in our object detection data set, with a consistent background. Future work will need to include a much larger data set of food items, methods for identifying previously unseen food items, and more sophisticated grasping and motion planning algorithms for handling complex food items in cluttered environments.

D. Implications and Future Directions

Based on the grasping performance of the system and the feedback gathered from the SCI participant in this study, we have identified three key areas of improvement:

1) Portability: The system should be entirely contained on a powered wheelchair. The robot and gripper should be more appropriately designed for feeding tasks. Ideally, the system should be scale-able to other ADLs as well. Recent work in this area can be found in [19].

2) Robust Object Detection and Grasp Planning: In a home environment, the system will undoubtedly encounter unknown objects. It must be capable of generalizing grasping strategies across various utensils and food items. Recent work in this area can be found in [20], [21].

3) Human-in-the-loop Adaptive Learning: As noted by our SCI participant, people will inevitably have different preferences and eating habits. By including the user in the learning process, they can train their own system to behave according to their specific needs, updating parameters such as robot speed, bite size, dwell time, bite transfer delay, etc. Recent work in this area can be found in [22], [23].

VI. Conclusion

In this work, we present a proof-of-concept demonstration for a robot-assisted feeding system for individuals with movement disorders. While the current task scope is limited, we consider the initial performance to be promising and justifies continued research in this area. We believe including the beneficiary community in the design process is critical to developing assistive technologies. We hope that by using feedback from individuals with SCI to drive design decisions, we may develop an RAF system that more closely aligns with the needs of people who have difficulty feeding themselves.

References

[1] W. H. Organization *et al.*, *World report on disability 2011*. World Health Organization, 2011.

[2] T. M. Gill, C. S. Williams, and M. E. Tinetti, "Assessing risk for the onset of functional dependence among older adults: the role of physical performance," *Journal of the American Geriatrics Society*, vol. 43, no. 6, pp. 603–609, 1995.

[3] V. Kumar, T. Rahman, and V. Krovi, "Assistive devices for people with motor disabilities," *Wiley Encyclopedia of Electrical and Electronics Engineering*, vol. 22, 1997.

[4] C. A. Stanger, C. Anglin, W. S. Harwin, and D. P. Romilly, "Devices for assisting manipulation: a summary of user task priorities," *IEEE Transactions on rehabilitation Engineering*, vol. 2, no. 4, pp. 256–265, 1994.

[5] (2021) Performance health. [Online]. Available: https://www.performancehealth.com/meal-buddy-systems#sin=44393

[6] (2021) Obi feeder. [Online]. Available: https://meetobi.com

[7] R. Feng, Y. Kim, G. Lee, E. K. Gordon, M. Schmittle, S. Kumar, T. Bhattacharjee, and S. S. Srinivasa, "Robot-assisted feeding: Generalizing skewering strategies across food items on a realistic plate," *arXiv preprint arXiv:1906.02350*, 2019.

[8] T. Bhattacharjee, G. Lee, H. Song, and S. S. Srinivasa, "Towards robotic feeding: Role of haptics in fork-based food manipulation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1485–1492, 2019.

[9] D. Park, Y. Hoshi, H. P. Mahajan, H. K. Kim, Z. Erickson, W. A. Rogers, and C. C. Kemp, "Active robot-assisted feeding with a general-purpose mobile manipulator: Design, evaluation, and lessons learned," *Robotics and Autonomous Systems*, vol. 124, p. 103344, 2020.

[10] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick, "Detectron2," https://github.com/facebookresearch/detectron2, 2019.

[11] D. Malyuta, "Guidance, Navigation, Control and Mission Logic for Quadrotor Full-cycle Autonomy," Master thesis, Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109, USA, Dec. 2017.

[12] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, 2004.

[13] I. Söderkvist and P.-Å. Wedin, "Determining the movements of the skeleton using well-configured markers," *Journal of biomechanics*, vol. 26, no. 12, pp. 1473–1477, 1993.

[14] D. Gallenberger, T. Bhattacharjee, Y. Kim, and S. S. Srinivasa, "Transfer depends on acquisition: Analyzing manipulation strategies for robotic feeding," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 267–276.

[15] (2021) Talon. [Online]. Available: https://talonvoice.com

[16] A. Bulat and G. Tzimiropoulos, "How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks)," in *International Conference on Computer Vision*, 2017.

[17] J. Brooke, "Sus: a "quick and dirty'usability," *Usability evaluation in industry*, vol. 189, no. 3, 1996.

[18] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Advances in psychology*. Elsevier, 1988, vol. 52, pp. 139–183.

[19] L. Birglen, "Design of a partially-coupled self-adaptive robotic finger optimized for collaborative robots," *Autonomous Robots*, vol. 43, no. 2, pp. 523–538, 2019.

[20] E. K. Gordon, X. Meng, T. Bhattacharjee, M. Barnes, and S. S. Srinivasa, "Adaptive robot-assisted feeding: An online learning framework for acquiring previously unseen food items," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9659–9666.

[21] A. ten Pas, M. Gualtieri, K. Saenko, and R. Platt, "Grasp pose detection in point clouds," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1455–1473, 2017.

[22] T. Bhattacharjee, M. E. Cabrera, A. Caspi, M. Cakmak, and S. S. Srinivasa, "A community-centered design framework for robot-assisted feeding systems," in *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*, 2019, pp. 482–494.

[23] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, "Continual lifelong learning with neural networks: A review," *Neural Networks*, vol. 113, pp. 54–71, 2019.