

Lessons Learned from Designing and Evaluating a Robot-assisted Feeding System for Out-of-lab Use

Amal Nanavati*, Ethan K. Gordon†, Taylor A. Kessler Faulkner*, Yuxin (Ray) Song*, Jonathan Ko§, Tyler Schrenk¶, Vy Nguyen‡, Bernie Hao Zhu*, Haya Bolotski||, Atharva Kashyap**, Sriram Kutty*, Raida Karim, Liander Rainbolt*, Rosario Scalise*, Hanjun Song||, Ramon Qu, Maya Cakmak* and Siddhartha S. Srinivasa*

*University of Washington, USA. {amaln, taylorkf, syx1995, haozhu, skutty, lrainb, rosario, mcakmak, siddh}@cs.uw.edu

†University of Pennsylvania, USA. ethankg@seas.upenn.edu ‡Hello Robot, USA. v.nguyen.ot@gmail.com

§ LinKo PLLC, USA. ¶ The Tyler Schrenk Foundation, USA

|| Massachusetts Institute of Technology, USA. ** University of Michigan, USA

Abstract—Millions of people cannot eat independently due to a disability, and caregiver-assisted meals can make them feel self-conscious, pressured, or burdensome. Robot-assisted feeding promises to empower people with motor impairments to feed themselves. However, current research typically examines specific robotic system subcomponents and evaluates them in controlled lab settings. This leaves a gap in developing and evaluating an end-to-end system that can feed entire meals in out-of-lab settings. We present one such system, which we developed collaboratively with two community researchers (CRs) with motor-impairments. The key challenge of developing a robot feeding system for out-of-lab use is the varied off-nominal scenarios that inevitably arise. Our key insight is that users can overcome many off-nominals, provided customizability and control over the system. Our system improves upon the state-of-the-art with: (1) a user interface that provides substantial user customizability and control, (2) a bite selection implementation that incorporates users-in-the-loop to generalize across food items, and (3) portable hardware that facilitates system use in diverse environments without inhibiting user mobility. We conduct two studies to evaluate the system. In Study 1, five users with motor impairments and one CR use the system to feed themselves meals of their choice in a cafeteria, office, or conference room. In Study 2, one CR uses the system in his home for five days, feeding himself 10 meals across diverse contexts. We present 3 key lessons learned: (1) spatial contexts are numerous, customizability lets users adapt to them; (2) off-nominals will arise, variable autonomy lets users overcome them; and (3) assistive robots’ benefits depend on context. We provide video footage and code on our website.¹

Index Terms—robot-assisted feeding; human-robot interaction

I. INTRODUCTION

Eating is a basic activity of daily living (ADL), one of the “fundamental skills required to independently care for oneself” [1]. Satisfaction with food-related matters is positively correlated with physical and mental health [2, 3, 4]. Unfortunately, for the millions of people who need caregiver assistance to eat,² mealtimes can lead to feelings of self-consciousness, pressure, and being burdensome to others. [6, 7].

Robot-assisted feeding is emerging as a promising way to alleviate these challenges [6, 7, 8]. Research in this area often

This work was (partially) funded by UW CREATE; UW Allen School Postdoc Research Award; NSF GRFP (DGE-1762114) and NRI (#2132848); DARPA RACER (#HR0011-21-C-0171); NIBIB (#1R01EB034580-01); and ONR (#N00014-24-S-B001, #2022-016-01 UW). We gratefully acknowledge gifts from Amazon, Collaborative Robotics, Cruise, and others.

¹<https://robotfeeding.io/publications/hri25a/>

²In 2010, 1.8M Americans needed assistance eating due to a disability [5].



Fig. 1. We evaluate the robot feeding system with: (Bottom) an $n = 5$ study across 3 out-of-lab locations; (Top) a 5-day, $n = 1$ in-home deployment.

focuses on specific technical components of eating, including bite acquisition [9, 10, 11, 12, 13], transfer [14, 15, 16], and timing [17, 18]. These contributions are evaluated via targeted studies that control for other aspects of eating, e.g., being in a controlled lab environment [14], limiting food positions and types [9, 19], and limiting the number of bites per user [15]. Such limitations are necessary to isolate the component under investigation from other meal-related factors. However, this leaves a gap in developing and evaluating an end-to-end system for robot-assisted feeding.

This paper addresses that gap. *Our goal is to develop an end-to-end robot feeding system that users can independently use to feed themselves meals of their choice outside the lab.* Except for system setup, onboarding, and pre-cut bite-sized meal preparation, users should be able to use the system to independently feed themselves entire meals.

The *key challenge* of developing a system for out-of-lab use is the wide variety of off-nominal scenarios that can arise: e.g., the user may cough; the robot may not acquire a food bite; or the plate may shift. Our *key insight* is that users can overcome many off-nominals, provided acceptable levels of

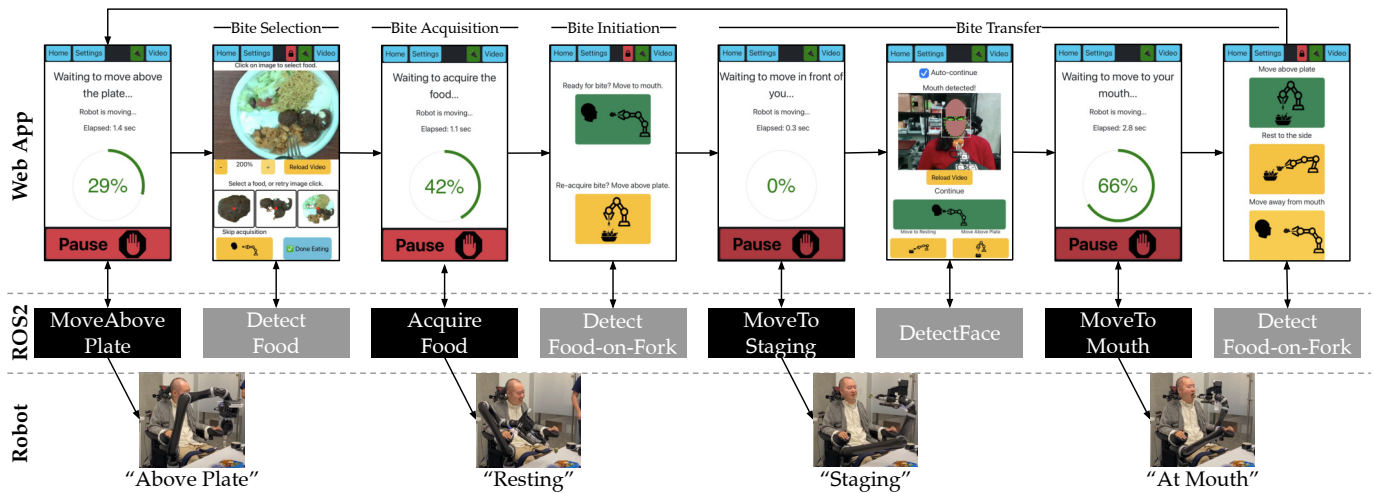


Fig. 2. An overview of the robot-assisted feeding system’s software. The user interacts with the system using the web app (top), which invokes ROS2 interfaces (middle) that move the robot to key configurations (black) and perceive the environment (grey). The default configurations are shown at the bottom.

customizability and control over the system.

We worked with two community researchers (CRs) with motor impairments to co-design and evaluate the feeding system. In Study 1, 5 participants and one CR³ used the robot to eat a meal of their choice in a cafeteria, office, or conference room. In Study 2, one CR used the robot *in his home* over 5 days to feed himself 10 meals in diverse contexts. Both studies were approved by our institution’s ethical review board.⁴

Our *key contribution* is a novel, open-source¹ robot-assisted feeding system (Fig. 2) that has demonstrated success feeding real meals to real users in real environments for over 15 hours. To our best knowledge, this has not been previously demonstrated by modern research system [15, 20, 21, 22]. A majority of users rated the system as being average-or-above in usability and outperforming caregivers in user independence and control. We attribute the system’s success to three key improvements over the state-of-the-art:

- 1) A unique *web app* that is the seat of system logic and provides substantial customizability and control to users.
- 2) A novel *bite selection implementation*, with user-in-the-loop input, to accommodate diverse foods.
- 3) *Portable, flexible hardware* (Fig. 3), facilitating system use in diverse environments without hindering user mobility.

We also contribute 3 *key lessons learned* from developing and deploying this system: (1) spatial contexts are numerous, customizability lets users adapt to them; (2) off-nominals will arise, variable autonomy lets users overcome them; and (3) assistive robots’ benefits depend on context.

II. RELATED WORK

History of Robot-assisted Feeding. Enabling people with motor impairments to eat independently has been a long-standing research goal [27, 28, 29]. Research in the 1970s included trained capuchin monkeys [30] and the Morewood Spoon Lifter, where a user shovels food into a pedal-controlled spoon using a head-mounted rod. This robot was sold as the




Fig. 3. The robot feeding system’s hardware, mounted on a wheelchair (left) or hospital table (right). Power and compute is self-contained; no wires leave.

“Winsford Feeder” [29] and was clinically evaluated in labs and homes [27]. In the 1980s, multi-purpose systems emerged that allowed feeding, brushing teeth, and more [31, 32].

Contemporary Robot-assisted Feeding. The last two decades have seen many commercial robotic feeding systems: Bestic [33], Obi [23], Neater Eater [24, 25], and more [28, 29, 34]. These table-mounted robots improved users’ mental and physical well-being [33, 35]; however, all but Obi and Neater Eater were discontinued. This may be due to an overreliance on fixed bite acquisition and transfer motions that led to acquisition failures, dropped food, or neck strain [26, 36, 37, 38].

Researchers address these limitations by integrating perception—e.g., cameras and force-torque (F/T) sensors—onto their robotic feeding systems. These robots can be mounted on wheelchairs [17, 19, 39] or tables (fixed [14] or portable [20]), or be mobile manipulators [21, 22].

Research on these systems often focuses on bite acquisition and bite transfer. *Bite acquisition* involves using a fork [10, 18, 40, 41], spoon [42, 43, 44], chopsticks [12, 45], or other tool [11, 21, 26] to grasp food. Utensils often follow motion primitives [10, 18, 39, 46] derivable from human data [9] and chained for complex motions [13, 41]. Online learning can improve primitive selection over time [40, 47]. *Bite transfer* involves handing off an acquired bite to the user’s mouth, using trajectories created with heuristic-based planning [14]

³All CRs and participants need caregiver assistance to eat.

⁴UW IRB: STUDY00005607, STUDY00020357

Robot	Approximate Cost	Mounting	Autonomous Motion	(General) Food Detection ⁵	Face Detection	Collision Detection / Avoidance	Portable & Self-contained	User Can Stop / Restart Motion	Customizable Robot Motion	Multiple UI Modalities
Obi [23]	\$8,625	Table	✓	(X) X	X	✓/ X	✓	✓/ ✓	✓	✓
Neater Eater [24, 25]	\$6,500	Table	✓	(X) X	X	X/ X	✓	✓/ ✓	✓	✓
Song et al. [20, 26]	–	Table	✓	(X) X	X	X/ X	✓	✓/ ✓	X	✓
Park et al. [21]	\$400,000	Mobile base	✓	(X) ✓	✓	✓/ ✓	X	✓/ ✓	✓	✓
Nguyen [22]	\$17,950	Mobile base	X	(X) X	✓	X/ X	X	✓/ ✓	X	✓
Bhattacharjee et al. [19]	\$50,000	Wheelchair	✓	(X) ✓	✓	✓/ ✓	X ⁶	✓/ X	X	✓
Jenamani et al. [13, 15]	\$50,000	Wheelchair	✓	(X) ✓	✓	✓/ ✓	X ⁶	✓/ X	X	X
This paper	\$50,000	Wheelchair or Table	✓	(✓) ✓	✓	✓/ ✓	✓	✓/ ✓	✓	✓

TABLE I

COMPARISON BETWEEN THIS AND OTHER ROBOT-ASSISTED FEEDING SYSTEMS (TOP: COMMERCIAL, BOTTOM: RESEARCH).

or learned from demonstrations [48]. Recent work studied in-mouth bite transfer for users without neck mobility [15, 16].

Other research includes: studying the coupling of acquisition and transfer [39], predicting food preference [49], detecting food [18, 50] or the mouth [51], predicting bite timing [17, 18], detecting anomalies [52], and studying natural language interfaces [53]. Simulation environments for robotic caregiving have also been developed [54, 55].

Table I compares technical capabilities across contemporary robot feeding systems. Ours differs from others in its general food detection capability and from wheelchair-mounted systems in its portability, customizability, and user control.

Out-of-lab Deployments. There is growing interest in out-of-lab deployments of physically assistive robots (PARs) [56]. This includes a robotic guide for blind museum visitors [57], a teleoperated mobile manipulator deployed over weeks that also fed its user [58, 59, 60], and a table-mounted robot-assisted feeding system [26]. These deployments, where users freely use systems in-the-wild, yield valuable insights about task nuances, user preferences, and needed system improvements.

III. SYSTEM

A. System Development

1) *Community-based participatory research (CBPR):* We conducted this research with two community researchers (CRs) following CBPR principles [61, 62]. CBPR maintains that community members and academic researchers have unique expertise and experiences, so addressing a community need requires sharing power, resources, and knowledge [63, 64]. CBPR is used in health sciences [62, 65] and increasingly in assistive technology (AT) research [6, 66, 67, 68, 69].

We met CR1 in 2018 through our network. He was passionate about assistive robots. “For a long time, I would only let my mom feed me. I wondered, why am I so uncomfortable with others feeding me that I’ll just not eat? I realized that eating is so individualized, with so many intricacies. If I can have a robot do it, I can learn to adapt to it, but it would be *me feeding me*, and that would be huge” (CR1). He participated in pilot studies and more, and in 2021 we began working with him as a CR.⁷ The first multi-day deployment was planned

⁵This refers to detecting food masks, irrespective of semantic labels.

⁶These systems have wires connecting the robot’s end-effector to external power or compute. This restricts robot motion and poses a trip hazard [15, 19].

⁷Semi-weekly meetings with CRs involved learning about their meal experiences, teaching them technical concepts, and iterating upon the system.

in his home, but he passed away months before. His friend, CR2, wanted to honor his legacy by continuing the work. We began working with CR2, culminating in a 5-day deployment in his home. Both CRs have quadriplegia due to a spinal cord injury (SCI) and are paper co-authors.

2) *An Analysis of Off-Nominals:* Past work surfaced the importance of anomaly monitoring, detection, and correction in robot-assisted feeding [7, 52, 70]. We worked with CR1 to compile a list of such off-nominal scenarios⁸ (Table II). Despite the diversity of off-nominals, CR1 observed that users could resolve many of them if provided control (e.g., to retry robot motions, to teleoperate the robot) and customizability (e.g., to adapt the robot to their environment).

3) *Guiding Design Principles:* Informed by formative research in robot-assisted feeding [6, 7, 8], we worked with CR1 to develop the following guiding design principles:

- **Portability.** The system must not hinder the user’s or others’ mobility. It must be easy to transport and set up.
- **Safety.** The system must not harm the user, other people, or objects in the environment.
- **Reliability.** The system must be able to reliably acquire and transfer the food items a user regularly eats.
- **Customizability.** The user should be able to customize the system to their contexts and preferences.
- **User Control.** The user should have fallback control and enough transparency into the system to utilize it.

B. System Overview

This section provides a system overview, including how we implement the “portability” and “safety” design principles. All software and custom hardware is available open-source.¹

1) *Hardware:* Fig. 3 shows system hardware.

Robot. A 6 degree-of-freedom (DoF) Kinova JACO arm.

Camera. An eye-in-hand system with an Intel Realsense D415 RGBD camera attached to an Nvidia Jetson Nano for wireless image transport. It accesses the robot’s internal power through a hole drilled above the last joint. Its position was designed to maintain continuous wrist rotation.

Fork. A custom 3D-printed fork assembly, held in the robot’s two-finger gripper. The fork has a 6 DoF ATI Nano25 F/T sensor attached to a battery-powered transmitter that charges with a magnetic connection to the eye-in-hand system.

⁸In an off-nominal scenario, something involved in system execution—the user, robot, or environment—does not proceed “according to plan” [71, 72].

User	Robot	Environment
User no longer wants the bite	Robot collides with object	Food falls off of the fork
User gets pulled into a conversation	Robot fails to perceive bite	Unexpected relative configuration of user/robot/plate
User cannot eat (e.g., is coughing)	Robot fails to acquire bite	Local area network fails
User takes a partial bite	Robot fails to perceive face	Device running the web app fails
User clicks an unintended button	Robot stops too far from face	Voice-based assistive technology fails (e.g., due to noise)

TABLE II

OFF-NOMINAL SCENARIOS THAT CAN ARISE DURING ROBOT-ASSISTED FEEDING, CO-CREATED WITH CR1.

Compute. A Lenovo Legion 5 laptop (RTX 3060 6GB GPU) that connects to the robot over USB and to a standard accessibility button for emergency stop (e-stop) over 3.5mm aux. The e-stop is mounted in user-accessible location.

Network. A local area network that enables system component communication. Users can either use a home router or the Cradlepoint IBR900 that travels with the system

Mount. A portable mount for the system. Components can be mounted to a wheelchair or hospital table (Fig. 3), or the robot can be on a tripod with other components in a backpack.

Power. A 24V DC power supply. This can be provided by a power wheelchair’s internal battery, a portable power station, or a wall outlet. For the former two, no wires leave the mount.

2) *Software:* Fig. 2 shows system software.

Hardware Interface and Controllers. The software stack is built on ROS2 and ros2-control. All controllers are “force-gated,” so execution is aborted if measured force or torque exceeds configurable thresholds.⁹ All Cartesian control uses a selectively damped pseudo-inverse Jacobian [74, 75].

Planning. We use MoveIt2 for planning, kinematics (`pick_ik` [76]), and collision detection. We use RRT-Connect [77] with shortcutting and hybridization [78] for planning due to its successful prior use [14]. The planning scene has a hull around the user and wheelchair, tight workspace walls, and an Octomap [79] for user-specific obstacles (e.g., ATs). We reject plans whose joint rotations exceed a threshold.

Robot Behaviors. The robot exposes modular behaviors through ROS2 actions and services. For example, “Segment-FromPoint” takes in a user-specified seed pixel in the robot’s image and returns contender masks of that food item. “AcquireFood” computes the food reference frame and executes an acquisition action. `py_trees` represents each motion-related action as a behavior tree (BT) [80], which encapsulates complex robot actions while re-using constituent behaviors.¹⁰

User Interface. Users interact with the robot via a React web app, accessed from any device with a browser. Thus, they can use *their own ATs* to interact with the system. The web app controls system execution, letting users navigate the state machine by invoking robot actions (Fig. 2). Unlike prior systems that have the robot control system execution [15, 19], our architecture increases robustness to off-nominals; users can at any time pause, go back, or redirect the system.

Safety Watchdog. A 60Hz watchdog verifies invariants, e.g., the F/T sensor and e-stop are connected and the e-stop has not been clicked. Robot motion stops if an “all clear”

⁹Thresholds are 1N when approaching the face, ≤ 50 N when acquiring foods, else 4N, far below max force standards for collaborative robots [73].

¹⁰We chose BTs over other models (e.g., finite state machines) for their readability and availability of documented open-source software [80, 81].

watchdog message is not received for 0.5 secs, simplifying verification of safety-critical code by centralizing it.

C. Robot-Assisted Feeding Procedure

This section focuses on the “reliability” design principle. Fig. 4 shows key components of the feeding procedure. Fig. 2 shows quoted robot arm configurations (below).

1) *Bite Selection:* Users specify their bite preference through a UI that was informed by a pilot study with CR1. At the “above plate” configuration, the robot sends the live RGB view to the web app via WebRTC. The user then selects a pixel on the image. `SegmentAnything` (Vit-B pre-trained) [82] generates 3 candidate masks, which are rendered on the web app with a dot showing roughly where the robot will skewer. The user can then select a candidate or re-select a pixel.

While waiting for bite selection, the system runs table detection to generalize across table heights. It uses OpenCV’s Hough Transform to detect the plate, takes the depth within a 50px ring around the plate, removes outliers, and fits a plane.

2) *Bite Acquisition:* The user-selected mask is sent to a policy that selects an acquisition action for the arm to execute.

Acquisition Action. This action is based on the specifications and schema defined in [9]. Each action consists of 3 linear Cartesian motions: *approach* (pre-contact), *grasp* (in-food manipulation), and *extraction*. The approach is defined by the initial fork orientation, the approach vector, and a target contact location on the food. The grasp and extraction are defined as a Cartesian twist (angular and linear velocity) and duration. Each motion has end-effector F/T thresholds that, when exceeded, abort the current motion and move to the next. All 3 motions are defined with respect to the food center (+x: the major axis of the bounding ellipse, -z: gravity).

We use 7 actions, based on those learned from human data in [9].¹¹ We manually adjust the actions to improve stability: we scale twists to a constant angular velocity, remove angular rotations $< 5^\circ$, and define the food frame to be the top instead of bottom of the bite (to align with the perceived depth).

Action Selection Policy. We implement [40]’s online learning system. The policy linearly maps the bite’s visual features (last layer of a custom-trained RetinaNet [39]) to the 7 primitives. Map parameters are learned online via LinUCB [83].

Post-Acquisition. The robot moves to a “resting” configuration. The user can initiate bite transfer when ready or have the robot move back “above plate” if acquisition failed.

3) *Bite Transfer:* The robot moves to a “staging” configuration with a view of the user’s face, detects it with the Haar Cascade classifier [84], uses Cartesian control to move the fork to the mouth, and then returns to “staging” and “above plate.”

¹¹Specifically, from [9], we use the 3 baseline actions, variants of those three with tilt-back extraction, and the human-informed action #3.

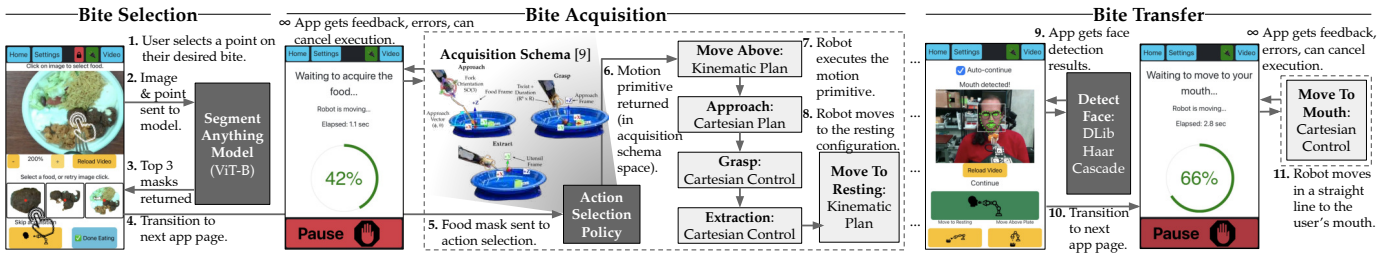


Fig. 4. A system diagram of bite selection, acquisition, and transfer, showing how the web app communicates with machine learned models (dark grey) and robot motion code (light grey). Components surrounded in a dashed line are represented as a BT. Acquisition schema visualization adapted from [9].

While by default the preceding transitions await user input, each has an optional auto-continue setting.

- 1) **Post-Acquisition.** The robot uses a food-on-fork detector to predict if acquisition succeeded and moves forward to “staging” configuration or back to “above plate.”
- 2) **Moving to Mouth.** The robot auto-continues once face detection perceives a face within the expected distance.
- 3) **Moving from Mouth.** The robot moves away if the food-on-fork detector (see Appendix) perceives no food.

D. Implementing the Remaining Design Principles

This section focuses on the remaining two design principles.

1) **Customizability:** Customizability is useful for ATs [85], PARs [56], and robot feeding [19, 86]. We provide it through:

Arm Configurations. Users have full control of the “above plate,” “resting,” and “staging” configurations, which are used as waypoints in all robot motions.

Bite Transfer. Users can customize how far from their mouth the robot stops and its speed near their mouth.

Auto-Continue. Users can customize whether the web app waits for their input or uses perception to transition states.

Planning Scenes. User can choose from pre-defined planning scenes: (1) the user and robot are on a wheelchair or (2) the user is in bed, and the robot is on a hospital table.

Customization is done via a web app settings menu. Given prior findings that users and caregivers tinker with assistive robots [38], we design the settings menus using the “Designing for Tinkerability” framework [87]. We provide “fluid experimentation” to users through direct access to the parameter space and “immediate feedback” by allowing them to try out the robot motions that result from their customizations.

2) **User Control:** Past work showed the importance of *variable autonomy* for PARs [58, 59, 88, 89]. Thus, we provide users multiple levels of control (LoCs), as defined in [90]. When the robot is moving, users have “supervisory control” to pause it. Doing so drops the LoC to “decision support,” where the robot provides multiple options for what to do next. At any time, users can drop the LoC to “teleoperation,” where the web app gives them direct Cartesian and joint control.

IV. STUDY 1: MULTI-USER, ON-CAMPUS STUDY

Study 1 quantitatively investigates: **How does the system perform across different users in out-of-lab settings?** To answer, we invited 5 participants and CR2 (Table III) to eat a meal of their choice in a campus cafeteria, conference room,

or office.¹² Following the CR insight that users focus on core AT features at first, we introduced users to all features but customizability and teleoperation.

After obtaining informed consent, we asked participants pre-meal questions while cutting the food, mounted the e-stop, loaded the web app on their device, and walked them through a bite. We explained system features but did not prescribe how participants should behave. Participants then began their meal. One researcher ate with them as a social partner, another took notes, and a third monitored system software to terminate code if necessary. After the meal, we asked post-meal questions. We conducted system patches between studies (see Appendix).

We collected data using evaluation indicators from [7], including *objective metrics* (e.g., meal time profile; acquisition and transfer success rate; system errors) and *subjective metrics* (pre/post ratings of caregiver and robot feeding, the NASA-TLX [91] for cognitive workload, and the System Usability Scale (SUS) [92]). As widely used metrics, the TLX and SUS have baselines from meta-analyses; the TLX’s is 37 ± 11 [91], and the SUS’s is a standardized grade where C is average [92].

A. Results

Table IV highlights results from Study 1.

1) **Bite Duration:** Bite duration¹³ ranged from 1:00–2:26 minutes.¹⁴ In contrast, people without disabilities take 18–30s [93, 94]. Fig. 5 shows the time profile for P4’s meal. Most time was spent in bite acquisition and moving to his mouth, which have the most interleaved perception, planning, and execution. The largest differences across participants depended on whether they used mouth-based (e.g., voice control, mouth joystick) or touch-based assistive technologies since users could not use the former while talking or chewing.

2) **Bite Acquisition and Transfer:** Prior work found that 80% acquisition success¹⁵ was sufficient for practical use [19]. For all users, the system neared or exceeded that for the most successful food items, and for all but two users it did so throughout the entire meal. The transfer success rate¹⁶ was

¹²Though not as controlled as labs, these are semi-controlled settings, e.g., with standardized lighting, less clutter, etc. Locations were chosen by availability (conference room, office) and user willingness (cafeteria).

¹³“Bite duration” excludes time between bites (e.g., conversations).

¹⁴System Patch 1 increased speed by 66%; P1 used the slower robot.


¹⁵A bite acquisition success is recorded if the food is on the fork at the end of acquisition; else, failure.

¹⁶A bite transfer success is recorded if the robot stops where the user can eat the bite; else, failure.

ID	Age	Gender	Impairment	Eating assistance providers	Feeds Self?	Study Device	Device interaction	Study location(s)	Selected meal(s) items ¹⁷
P1	49	M	C3 SCI ¹⁸	Parent(s)	Never	Phone	Voice control	Conference room	Pizza, broccoli
P2	42	F	C5 SCI	FCs, parent(s)	Never	Phone	Stylus	Office	Chicken, salad
P3	45	M	Arthrogyposis	FCs, spouse	Sometimes	Phone	Stylus	Conference room	Sandwich, brownies
P4	62	M	C3 SCI	FCs	Never	Phone	Touch	Office	Chicken, potatoes
P5	61	F	C5-6 SCI	FCs, spouse	Sometimes	Tablet	Touch	Office	Salmon, brussels
CR2	43	M	C2 SCI	FCs	Never	Phone	Mouth joystick	Cafeteria	Stir-fry beef, tofu

TABLE III

PARTICIPANT DEMOGRAPHICS AND DETAILS FOR STUDY 1. FC REFERS TO FORMAL CAREGIVERS, I.E., PAID AND TRAINED PROFESSIONALS.

Fig. 5. How long each stage of feeding took across P4’s 30 successful bites¹⁹.

94% for P1 and 100% for the others.

3) *Off-nominal Scenarios*: Each meal had off-nominal scenarios. Many were user recoverable (e.g., 16% for P1, 88% for P2) via the web app; these included acquisition and transfer failures, robot action errors, mistaken app clicks, and browser interruptions. For some off-nominals, researchers intervened physically (e.g., moving the plate, re-aligning the fork in the gripper) or digitally (e.g., restarting code).

4) *Caregiver Feeding Comparison*: Fig. 6 shows user ratings of caregiver vs. robot feeding. Robot feeding outperforms in users’ sense of control (Q1-2) and independence (Q3). 4/5 participants and CR2 agreed that “When I ate with the robot, I was confident that I would remain safe” (Q6).

5) *Cognitive Workload*: All users but P3 reported experiencing a cognitive workload below the baseline. This indicates that the cognitive workload required to use the system was relatively low despite its many user-in-the-loop components.

6) *System Usability*: Three of five participants and CR2 rated the system as average-or-above average usability. There was wide variability, from P3’s F to P4’s A+. This variability is to be expected given the diversity of users. For example, P3’s current self-feeding technique provides many of the robot’s functionalities, but he imagined it would help “others who can’t use a [self-feeding] system like me.”

V. STUDY 2: SINGLE-USER, IN-HOME DEPLOYMENT

Study 2 qualitatively investigates: **How does the system perform across the diverse contexts that arise when eating in the home?** To answer, we deployed the robot in CR2’s home for 5 consecutive days to help him eat 2 meals/day. Pre-deployment, we worked with CR2 and an occupational therapist (OT) to identify CR2’s meal-related contexts and goals. Context is “any information that can...characterize the situation of entities...considered relevant to the interaction between a user and an application” [96]. These included:

- **Spatial Context**. CR2 cannot sit up for consecutive days and so alternates between bed and wheelchair days.

¹⁷Meals were bought from restaurants or made by the user’s caregivers.

¹⁸SCI severity is classified by the injured vertebra; C1 is nearest the neck.

¹⁹Box: 25th, 50th, 75th percentile. Whiskers: 1.5-IQR. Outliers excluded.




Fig. 6. Users self-reported comparison: eating with caregivers vs. the robot.

- **Social Context**. CR2 has three caregivers, C1-C3 (see Appendix), who typically feed him.
- **Temporal Context**. Mornings are busy with CR2’s care routine and daytimes with work, but evenings are relaxed.
- **Activity Context**. CR2’s deployment goals were to (1) feed himself dinner while watching television, (2) spend time with a caregiver while both eat dinner, (3) feed himself while a caregiver does other care work, (4) feed himself breakfast while working, and (5) feed himself a mid-day snack while working.
- **Food Context**. CR2 is a flexible eater, enjoying ramen, pizza, chicken teriyaki, fruits/vegetables, and more.

On Mon, Wed, and Fri (wheelchair days), CR2 used the robot for breakfast and dinner. On Tues and Thurs (bed days) he used it for snack and dinner. CR2 selected meal locations and times. Before meals, one researcher set up the robot and acquired test bites, while another cut the food. We then brought the robot to CR2 and positioned the e-stop near him. He then customized the system and began his meal. One researcher monitored software; the other 1–2 took notes. After each meal, we had a semi-structured interview with CR2 and/or his caregivers; all gave informed consent before the study.

A. Results and Lessons Learned

CR2 used the robot to fed himself all 10 meals, including store-bought foods (e.g., fruits) and caregiver-prepared ones (e.g., avocado toast); he ate various cuisines (e.g., pizza, chicken teriyaki, charcuterie). Using the Medicare Section GG scale [97], the OT assessed that due to system use, CR2’s *level of independence during meals increased* from “dependent” (his baseline) to “supervision,” where the caregiver is on standby to provide intermittent assistance. We now present qualitative findings grouped into key lessons learned.

1) *Spatial Contexts are Numerous, Customizability Lets Users Adapt to Them*: The home setting’s spatial contexts differed from campus settings. There were *many environmental objects*: CR2 had a mouth joystick near his face and a laptop or phone in front, often on a face-height hospital table. These objects and the e-stop constrained the robot’s motion enough that CR2 sometimes said it was “threading a needle.” *Spatial*

ID	Meal Time	Bites Eaten	Median Bite Time (IQR)	User-resolved Off-nominals	Researcher Interventions (Physical, Software)	Acquisition Success Rate	Most Successful Food	Cognitive Workload (Baseline: 37 [95])	Usability Grade (Baseline: C [92])
P1	52:37	15	2:26 (0:54)	8	1, 0	0.79 (13/19)	Pizza: 0.78 (14/18)	17.50	D
P2	54:53	24	1:10 (0:12)	2	5, 5	0.65 (24/37)	Chicken: 0.85 (11/13)	29.17	C
P3	54:06	31	1:00 (0:09)	7	6, 1	0.69 (31/45)	Sandwich: 0.94 (14/17)	38.33	F
P4	56:52	30	1:10 (0:21)	22 ²⁰	2, 1	0.88 (30/34)	Chicken: 1.0 (13/13)	20.00	A+
P5	51:05	23	1:15 (0:20)	5	0, 1	0.79 (23/29)	Brussels: 0.86 (9/7)	19.17	B+
CR2	28:44	14	1:41 (0:24)	3	1, 1	0.78 (14/18)	Tofu: 1.0 (3/3)	19.17	A

TABLE IV

STUDY 1: PER-PARTICIPANT TIME PROFILE (MINS:SECS), NUMBER OF INTERVENTIONS, ACQUISITION RESULTS, AND SUBJECTIVE RESULTS.

configurations varied between the robot, user, and plate: the bed’s tilt, height, and the user’s lateral position varied; on a wheelchair, the plate’s position, height, and chair’s tilt varied. *Lighting conditions varied*: sources of light included windows, lamps, and ceiling lights, and many surfaces were white or reflective, creating backlight, reflections, and shadows.

To enable the robot to work given these varied spatial contexts, CR2 started all meals by checking the previously customized configurations relative to the current meal’s context. First, he customized the configurations to account for context, e.g., changing the “above plate” configuration to be above the plate. Second, he customized for preferences, e.g., trying to adjust “staging” to approach him below his eyeline.

Transparency into the downstream impacts of changes was crucial to this process. CR2 iteratively tuned the “resting” configuration and tried motions to/from it until finding one that gave his computer a wide berth. He iteratively tuned “staging” and checked face detection’s precision until he found a configuration with reliable face detection in that context.

This process also involved environmental modifications. CR2 sometimes asked a caregiver or us to adjust his laptop or mouth joystick to give the robot arm more room. Once, after realizing that forehead reflections were causing false positive face detections, he had a caregiver place a cap on his head.

Having access to the right level of customization was vital: “This was the sweet spot. I don’t want to have to type in code.” As we demonstrated the planning scene for CR2 and discussed ways for him to customize it, he said, “Totally automate that. Just the thought of it makes my head hurt.”

These observations led to the following lesson learned. *Tinkering is vital for assistive robots to work in users’ contexts [38]. Systems should be customizable to foster ease of tinkering. This requires intuitive control over parameters and transparency into the downstream impacts of parameters.*

2) *Off-nominals Will Arise, Variable Autonomy Lets Users Overcome Them*: Although customizability enabled the user to adapt the robot to spatial contexts, some contexts remained challenging for autonomous behaviors. First, some lighting conditions lowered face detection’s precision so much that CR2 did not want the robot to autonomously approach him: “[I don’t want it] nearing my eyes.” Second, some spatial configurations between the robot and plate hindered bite acquisition. At times, this was due to planning failures: all plans to move above the bite were rejected due to a joint rotation that exceeded thresholds. At other times, this was due

to off-centering: a camera enclosure screw hole was damaged in transit, changing the camera’s extrinsics.²¹ In both cases, variable autonomy let CR2 overcome these off-nominals.

To address face detection failures, if customization did not work CR2 bypassed face detection altogether by teleoperating the robot from “resting” to his mouth. This could be mentally taxing due to many Cartesian and joint motions. Thus, CR2 devised a novel way to customize the “resting” configuration so a single joint-1 rotation moved the fork directly to his mouth, reducing teleoperation to a single button press. Thus, *variable autonomy helped him overcome the off-nominal, but customization let him lower the cognitive workload involved.*

To address bite acquisition failures, CR2 interspersed teleoperation with autonomous robot behaviors. He occasionally teleoperated before autonomous acquisition, changing the robot’s starting configuration to overcome planning failures. At other times he used autonomous acquisition to move the fork above the food and then paused it, teleoperating the remainder. Yet other times he teleoperated the entire acquisition. These multiple levels of autonomy helped him avoid frustration: “I would [get frustrated] if it wasn’t working, and it just kept on doing it and doing it. I’d be like, ‘Oh, stop. Just give me regular control.’ But [with this system], it is within my control.”

All-in-all, for 9/10 meals the robot successfully autonomously transferred $\geq 80\%$ of bites; for 5/10 meals it successfully autonomously acquired $\geq 80\%$ of bites. Other meals combined autonomy with teleoperation. Despite needing to occasionally teleoperate, CR2 still found the robot empowering to use. “Sometimes people feed me, and I don’t like how they’re doing it. It’s weirdly empowering, as someone who’s been paralyzed as long as I have, to say, ‘I’m going to eat this. It’ll take me 3 times as long, but I’m not going to be frustrated while I eat.’”

These observations led to the following lesson learned. *The challenging contexts present in home environments can hinder a robot’s autonomous behaviors. Users can help the robot navigate through these challenges provided varied ways to control it. The robot’s benefit may outweigh the users’ cognitive workload required to control it when autonomy fails.*

The above two lessons validate our key insight about the importance of customizability and control. During system development, we under-appreciated *how often* customizability or control would be needed. However, because they were provided, CR2 leveraged them to resolve scenarios we had not anticipated and to develop novel strategies for system use.

²⁰Full-screen pop-ups appeared often, hindering P4’s web app use.

²¹Plate-to-camera distances varied ≥ 10 cm across meals. So even a 0.05rad extrinsics error could move the fork 5mm off-center, making foods roll away.

3) *Assistive Robots’ Benefits Depend on Context*: Contexts beyond the spatial affected CR2’s meal experiences.

Activity Context. CR2 attempted to perform each of his 5 goals at least once during the deployment. He felt he achieved the first 3. “Eating while watching TV; it’s totally possible. [Using the robot] is not distracting to the point where you can’t do it. I’ve also accomplished eating while my caregiver did something else, because C1 did laundry.” He also ate dinner alongside his caregiver (Fig. 1). However, CR2 could not achieve his latter 2 goals. “[I couldn’t] eat while working [because of] my over-expectation to not pay attention to the robot. And I had to pay attention to it.” One reason is overlapping demands on his faculties: CR2 types using dictation and so cannot type while chewing; he reads visually and so cannot do it while looking at the robot.²² CR2’s perspective on these goals shifted over the deployment. “I realized, ‘Food is important. You need to eat more than you need to finish work. And doing that is worth your attention.’”

Social Context. When present, caregivers participated in the meals. “I was involved, but he was doing everything by himself. I was [checking if] the robot dropped something, [giving him] a napkin, refilling the plate. When he chewed, I watched the movie.” (C1). Caregivers were also involved in food preparation. After seeing the robot acquire her carrots but not her zucchini, C3 said, “[Today] it was [too] soft or small. That could be improved if I prepared [meals] several times with the robot.” Importantly, CR2’s envisioned future robot use was conditioned on caregiver effort: “[If it takes too long], they will say, ‘Let me just feed you and not set up [and tinker with] the robot.’ And that would be reasonable.”

Researchers were also part of the social context. “If a plate were there the whole day—if you guys weren’t here—I would’ve gotten the work done. I would’ve taken one bite, waited 30 minutes, finished a task, and taken another bite. I would do it guilt-free [because no one has to wait for me].”

Food Context. Caregivers had concerns about the robot’s acquisition limits. “We need to choose foods that CR2 likes and the robot can pick up.” (C3). “[The robot currently] has too many limits with his diet and what he likes” (C1). CR2 envisioned food-dependent robot use. “I wouldn’t eat all my meals with it. Some foods I like [e.g., ramen] can be difficult for it. [But] I like pizza a lot; it did fine with pizza.”

Temporal Context. CR2 enjoyed the robot more during dinner. “[When I’m] eating for enjoyment, during dinner, [using the robot] is great. For breakfast and snack, where I feel I should be working, things are rushed.” His caregivers agreed: “Sometimes, CR2’s in rush. So we don’t have time to set [the robot] up. So we have to feed him.” (C2). Despite contextual differences, CR2 found the robot rewarding to use. “That Wednesday morning, there was a flow state. I was succeeding at such a rate that it felt good. I was like, ‘We’re getting into it, no matter how long it takes.’ At that point, my satisfaction levels are really high.”

These observations led to the following lesson learned. *Assistive robots integrate into a user’s life. They provide*

benefits in some contexts but not others. Such contextual benefits may still be sufficient to make them a valuable addition to the tools users and caregivers adopt for ADLs.

VI. LIMITATIONS AND FUTURE WORK

In this work, we collaborated with community researchers (CRs) to develop a robot-assisted feeding system that people with motor impairments can use to independently feed themselves outside of lab settings. We evaluated the system quantitatively with 5 users and CR2 in 3 locations (Sec. IV) and qualitatively in CR2’s home for 5 days (Sec. V). Although this work made progress towards our goal (Sec. I), results reveal system limitations to address to fully reach the goal.

For **bite acquisition**, a limitation was missed bites. An important future direction is incorporating closed-loop feedback into action primitives, e.g., adjusting the motion if the bite starts tilting or the fork fails to pierce it. Another important direction is expanding the food types the robot can acquire to include e.g., ramen. For **bite transfer**, a limitation was that no matter how much customization CR2 tried, the robot approached at his eyeline; future work should orchestrate transfer motions to approach from below. For **customizability**, users must be able to customize the planning scene; in our system, users had to choose among hard-coded scenes. For **user control**, the system must provide ways for users and caregivers to debug system problems, such as transparently explaining what error the robot encountered and how they could resolve it. For **user comfort**, approaches like compliant control from physical human-robot-interaction could improve the system [15, 98, 99]. For **commercial viability**, future work should focus on reducing system cost. Finally, co-designing setup and maintenance procedures with caregivers could improve **integration into care routines**.

This in-home deployment is only the beginning. An open-source, deployable system lays the foundation for: (1) follow-up deployments with CR2 as the system matures and (2) in-home deployments with other users to study additional meal contexts and address potential CR biases. Long-term, researchers should not be present since our presence influences system use (Sec. V-A3).

VII. SUPPLEMENTARY MATERIALS

Supplementary materials, hosted on Open Science Foundation at [100], include per-meal event annotations, user quotes, and appendices with system and study details.

ACKNOWLEDGMENT

This paper is dedicated to the memory of Tyler Schrenk (1985–2023; www.thetsf.org), a dear collaborator, advocate, and friend. We are deeply grateful to Tyler (CR1) and Jonathan Ko (CR2) for their involvement throughout this work. We acknowledge Daniel Gallenberger, Tapomayukh Bhattacharjee, and the entire Personal Robotics Lab for contributions to the initial version of the system. We thank Sandy Kaplan and the anonymous reviewers for their feedback.

²²In contrast, CR2 can look at the robot while *listening to a TV show*.

REFERENCES

- [1] P. F. Edemekong, D. Bomgaars, S. Sukumaran, and S. B. Levy, "Activities of daily living," 2019.
- [2] K. G. Grunert, M. Dean, M. M. Raats, N. A. Nielsen, M. Lumbers *et al.*, "A measure of satisfaction with food-related life," *Appetite*, vol. 49, no. 2, pp. 486–493, 2007.
- [3] J. Utter, S. Denny, R. Peiris-John, E. Moselen, B. Dyson, and T. Clark, "Family meals and adolescent emotional well-being: findings from a national study," *Journal of nutrition education and behavior*, vol. 49, no. 1, pp. 67–72, 2017.
- [4] K. Charras and M. Frémontier, "Sharing meals with institutionalized people with dementia: a natural experiment," *Journal of Gerontological Social Work*, vol. 53, no. 5, pp. 436–448, 2010.
- [5] K. A. Theis, A. Steinweg, C. G. Helmick, E. Courtney-Long, J. A. Bolen, and R. Lee, "Which one? what kind? how many? types, causes, and prevalence of disability among us adults," *Disability and health journal*, vol. 12, no. 3, pp. 411–421, 2019.
- [6] A. Nanavati, P. Alves-Oliveira, T. Schrenk, E. K. Gordon, M. Cakmak, and S. S. Srinivasa, "Design Principles for Robot-Assisted Feeding in Social Contexts," in *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, ser. HRI '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 24–33.
- [7] T. Bhattacharjee, M. E. Cabrera, A. Caspi, M. Cakmak, and S. S. Srinivasa, "A community-centered design framework for robot-assisted feeding systems," in *International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)*. ACM, 2019, pp. 482–494.
- [8] M. Pascher, A. Baumeister, S. Schneegass, B. Klein, and J. Gerken, "Recommendations for the development of a robotic drinking and eating aid—an ethnographic study," in *Human-Computer Interaction—INTERACT 2021: 18th IFIP TC 13 International Conference, Bari, Italy, August 30–September 3, 2021, Proceedings, Part I 18*. Springer, 2021, pp. 331–351.
- [9] E. K. Gordon, A. Nanavati, R. Challa, B. H. Zhu, T. A. K. Faulkner, and S. Srinivasa, "Towards general single-utensil food acquisition with human-informed actions," in *Conference on Robot Learning*. PMLR, 2023, pp. 2414–2428.
- [10] P. Sundaresan, J. Wu, and D. Sadigh, "Learning sequential acquisition policies for robot-assisted feeding," *arXiv preprint arXiv:2309.05197*, 2023.
- [11] J. Grannen, Y. Wu, S. Belkhale, and D. Sadigh, "Learning bimanual scooping policies for food acquisition," *arXiv preprint arXiv:2211.14652*, 2022.
- [12] T. Oka, J. Solis, A.-L. Lindborg, D. Matsuura, Y. Sugahara, and Y. Takeda, "Kineto-elasto-static design of underactuated chopstick-type gripper mechanism for meal-assistance robot," *Robotics*, vol. 9, no. 3, p. 50, 2020.
- [13] R. K. Jenamani, P. Sundaresan, M. Sakr, T. Bhattacharjee, and D. Sadigh, "Flair: Feeding via long-horizon acquisition of realistic dishes," *arXiv preprint arXiv:2407.07561*, 2024.
- [14] S. Belkhale, E. K. Gordon, Y. Chen, S. Srinivasa, T. Bhattacharjee, and D. Sadigh, "Balancing efficiency and comfort in robot-assisted bite transfer," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 4757–4763.
- [15] R. K. Jenamani, D. Stabile, Z. Liu, A. Anwar, K. Dimitropoulou, and T. Bhattacharjee, "Feel the bite: Robot-assisted inside-mouth bite transfer using robust mouth perception and physical interaction-aware control," in *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, 2024, pp. 313–322.
- [16] L. Shaikewitz, Y. Wu, S. Belkhale, J. Grannen, P. Sundaresan, and D. Sadigh, "In-mouth robotic bite transfer with visual and haptic sensing," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9885–9895.
- [17] J. Ondras, A. Anwar, T. Wu, F. Bu, M. Jung, J. J. Ortiz, and T. Bhattacharjee, "Human-robot commensality: Bite timing prediction for robot-assisted feeding in groups," in *Conference on Robot Learning*. PMLR, 2022.
- [18] L. V. Herlant, "Algorithms, implementation, and studies on eating with a shared control robot arm," Ph.D. dissertation, Carnegie Mellon University, 2018.
- [19] T. Bhattacharjee, E. K. Gordon, R. Scalise, M. E. Cabrera, A. Caspi, M. Cakmak, and S. S. Srinivasa, "Is more autonomy always better? exploring preferences of users with mobility impairments in robot-assisted feeding," in *2020 15th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2020, pp. 181–190.
- [20] W.-K. Song and J. Kim, "Novel assistive robot for self-feeding," *Robotic Systems—Applications, Control and Programming*, pp. 43–60, 2012.
- [21] 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.
- [22] V. Nguyen, "Increasing independence with stretch: A mobile robot enabling functional performance in daily activities," 2021.
- [23] D. Artman, M. Parkinson, and P. Adanti-Joy, "New Obi Robotic Dining Device a breakthrough for people living with physical challenges," Jul 2016.
- [24] J. Michaelis, "Mechanical methods of controlling ataxia." *Bailliere's Clinical Neurology*, vol. 2, no. 1, pp. 121–139, 1993.
- [25] M. Hillman, "Rehabilitation robotics from past to present—a historical perspective," *Advances in rehabilitation robotics: human-friendly technologies on movement assistance and restoration for people with disabilities*, pp. 25–44, 2004.
- [26] W.-K. Song, W.-J. Song, Y. Kim, and J. Kim, "Usability test of knrc self-feeding robot," in *2013 IEEE 13th*

International Conference on Rehabilitation Robotics (ICORR). IEEE, 2013, pp. 1–5.

- [27] G. N. Phillips, “Feasibility study for assistive feeder,” *Southwest Research Institute*, p. 1–130, 1987. [Online]. Available: <https://robotfeeding.io/publications/archive/phillips1987feasibility.pdf>
- [28] “Assistive dining technology 1970 through 2005,” *Mealtime Partners, Inc.*, p. 1–37, 2005. [Online]. Available: <https://robotfeeding.io/publications/archive/mealtime2005assistive.pdf>
- [29] I. Naotunna, C. J. Perera, C. Sandaruwan, R. Gopura, and T. D. Lalitharatne, “Meal assistance robots: A review on current status, challenges and future directions,” in *2015 IEEE/SICE international symposium on system integration (SII)*. IEEE, 2015, pp. 211–216.
- [30] J. T. MacFayden, “Educated monkeys help the disabled to help themselves,” *Smithsonian*, vol. 17, pp. 125–132, 1986.
- [31] M. Topping, “An overview of the development of handy 1, a rehabilitation robot to assist the severely disabled,” *Artificial Life and Robotics*, vol. 4, no. 4, pp. 188–192, 2000.
- [32] W. Seamone and G. Schmeisser, “Early clinical evaluation of a robot arm/worktable system for spinal-cord-injured persons,” *Journal of rehabilitation research and development*, vol. 22, no. 1, pp. 38–57, 1985.
- [33] A.-L. Lindborg and M. Lindén, “Development of an eating aid—from the user needs to a product,” in *pHealth 2015*. IOS Press, 2015, pp. 191–198.
- [34] R. Soyama, S. Ishii, and A. Fukase, “The development of meal-assistance robot ‘my spoon’,” in *Proceedings of the 8th International Conference on Rehabilitation Robotics, 2003*, 2003, pp. 88–91.
- [35] A. Kumar and M. Frances Phillips, “Use of powered mobile arm supports by people with neuromuscular conditions,” *Journal of Rehabilitation Research & Development*, vol. 50, no. 1, 2013.
- [36] J. Kim, W.-K. Song, S. Han, J. Heo, and K. Lee, “A study on the scooping and serving korean food using meal assistance robot,” *Gerontechnology*, vol. 21, 2022.
- [37] S. Ljungblad, “Applying “designerly framing” to understand assisted feeding as social aesthetic bodily experiences,” *ACM Transactions on Human-Robot Interaction*, vol. 12, no. 2, pp. 1–23, 2023.
- [38] N. C. M. Nickelsen, “Imagining and tinkering with assistive robotics in care for the disabled,” *Paladyn, Journal of Behavioral Robotics*, vol. 10, no. 1, pp. 128–139, 2019.
- [39] 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.
- [40] 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)*, 2020, pp. 9659–9666.
- [41] P. Sundaresan, J. Wu, and D. Sadigh, “Learning sequential acquisition policies for robot-assisted feeding,” *arXiv preprint arXiv:2309.05197*, 2023.
- [42] T. Rhodes and M. Veloso, “Robot-driven trajectory improvement for feeding tasks,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 2991–2996.
- [43] Y. Ohshima, Y. Kobayashi, T. Kaneko, A. Yamashita, and H. Asama, “Meal support system with spoon using laser range finder and manipulator,” in *2013 IEEE workshop on robot vision (WORV)*. IEEE, 2013, pp. 82–87.
- [44] S. Ishii, S. Tanaka, and F. Hiramatsu, “Meal assistance robot for severely handicapped people,” in *Proceedings of 1995 IEEE International Conference on Robotics and Automation*, vol. 2. IEEE, 1995, pp. 1308–1313.
- [45] A. Yamazaki and R. Masuda, “Autonomous foods handling by chopsticks for meal assistant robot,” in *ROBOTIK 2012; 7th German Conference on Robotics*. VDE, 2012, pp. 1–6.
- [46] 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*, 2019.
- [47] E. K. Gordon, S. Roychowdhury, T. Bhattacharjee, K. Jamieson, and S. S. Srinivasa, “Leveraging post hoc context for faster learning in bandit settings with applications in robot-assisted feeding,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 10 528–10 535.
- [48] G. Canal, G. Alenyà, and C. Torras, “Personalization framework for adaptive robotic feeding assistance,” in *International conference on social robotics*. Springer, 2016, pp. 22–31.
- [49] S. Javdani, H. Admoni, S. Pellegrinelli, S. S. Srinivasa, and J. A. Bagnell, “Shared autonomy via hindsight optimization for teleoperation and teaming,” *The International Journal of Robotics Research*, vol. 37, no. 7, pp. 717–742, 2018.
- [50] R. Feng, Y. Kim, G. Lee, E. K. Gordon, M. Schmitte, S. Kumar, T. Bhattacharjee, and S. S. Srinivasa, “Robot-assisted feeding: Generalizing skewering strategies across food items on a plate,” in *The International Symposium of Robotics Research*. Springer, 2019, pp. 427–442.
- [51] C. J. Perera, T. D. Lalitharatne, and K. Kiguchi, “Eeg-controlled meal assistance robot with camera-based automatic mouth position tracking and mouth open detection,” in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 1760–1765.
- [52] D. Park, H. Kim, Y. Hoshi, Z. Erickson, A. Kapusta, and C. C. Kemp, “A multimodal execution monitor with anomaly classification for robot-assisted feeding,” in *2017 IEEE/RSJ International Conference on Intelligent*

- Robots and Systems (IROS)*. IEEE, 2017, pp. 5406–5413.
- [53] A. Padmanabha, J. Yuan, J. Gupta, Z. Karachiwalla, C. Majidi, H. Admoni, and Z. Erickson, “Voicepilot: Harnessing llms as speech interfaces for physically assistive robots,” *arXiv preprint arXiv:2404.04066*, 2024.
- [54] R. Ye, W. Xu, H. Fu, R. K. Jenamani, V. Nguyen, C. Lu, K. Dimitropoulou, and T. Bhattacharjee, “Rcare world: A human-centric simulation world for caregiving robots,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 33–40.
- [55] R. Madan, R. K. Jenamani, V. T. Nguyen, A. Moustafa, X. Hu, K. Dimitropoulou, and T. Bhattacharjee, “Spars: Structuring physically assistive robotics for caregiving with stakeholders-in-the-loop,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 641–648.
- [56] A. Nanavati, V. Ranganeni, and M. Cakmak, “Physically assistive robots: A systematic review of mobile and manipulator robots that physically assist people with disabilities,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, 2023.
- [57] S. Kayukawa, D. Sato, M. Murata, T. Ishihara, H. Takagi, S. Morishima, and C. Asakawa, “Enhancing blind visitor’s autonomy in a science museum using an autonomous navigation robot,” in *Conference on Human Factors in Computing Systems (CHI)*. ACM, 2023, pp. 1–14.
- [58] A. Padmanabha, J. Gupta, C. Chen, J. Yang, V. Nguyen, D. J. Weber, C. Majidi, and Z. Erickson, “Independence in the home: A wearable interface for a person with quadriplegia to teleoperate a mobile manipulator,” in *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, 2024, pp. 542–551.
- [59] V. Ranganeni, V. Nguyen, H. Evans, J. Evans, J. Mehu, S. Olatunji, W. Rogers, A. Edsinger, C. Kemp, and M. Cakmak, “Robots for humanity: In-home deployment of stretch re2,” in *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, 2024, pp. 1299–1301.
- [60] S. A. Olatunji, V. Nguyen, M. Cakmak, A. Edsinger, C. C. Kemp, W. A. Rogers, and H. P. Mahajan, “Immersive participatory design of assistive robots to support older adults,” *Ergonomics*, pp. 1–15, 2024.
- [61] K. Hacker, *Community-based participatory research*. Sage publications, 2013.
- [62] M. Viswanathan, A. Ammerman, E. Eng, G. Garlehner, K. N. Lohr, D. Griffith, S. Rhodes, C. Samuel-Hodge, S. Maty, L. Lux *et al.*, “Community-based participatory research: Assessing the evidence: Summary,” *AHRQ evidence report summaries*, 2004.
- [63] K. J. Peralta, “Politics of knowledge in community-based work,” in *Dimensions of Community-Based Projects in Health Care*. Springer, 2018, pp. 67–78.
- [64] M. Minkler and N. Wallerstein, *Community-based participatory research for health: From process to outcomes*. John Wiley & Sons, 2011.
- [65] B. A. Israel, A. J. Schulz, E. A. Parker, and A. B. Becker, “Review of community-based research: assessing partnership approaches to improve public health,” *Annual review of public health*, vol. 19, no. 1, pp. 173–202, 1998.
- [66] J. Mankoff, G. R. Hayes, and D. Kasnitz, “Disability studies as a source of critical inquiry for the field of assistive technology,” in *International ACM SIGACCESS conference on Computers and accessibility (ASSETS)*. ACM, 2010, pp. 3–10.
- [67] D. Chapko, P. Frumiento, N. Edwards, L. Emeh, D. Kennedy, D. McNicholas, M. Overton, M. Snead, R. Steward, J. M. Sutton *et al.*, ““we have been magnified for years-now you are under the microscope!”: Co-researchers with learning disabilities created an online survey to challenge public understanding of learning disabilities,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–17.
- [68] R. S. Kushalnagar and C. Vogler, “Teleconference accessibility and guidelines for deaf and hard of hearing users,” in *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, 2020, pp. 1–6.
- [69] J. Bell and T. W. Leong, “Collaborative futures: Co-designing research methods for younger people living with dementia,” in *Proceedings of the 2019 chi conference on human factors in computing systems*, 2019, pp. 1–13.
- [70] A. Nanavati, P. Alves-Oliveira, T. Schrenk, E. K. Gordon, M. Cakmak, and S. S. Srinivasa, “Unintended failures of robot-assisted feeding in social contexts,” in *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, 2023, pp. 884–886.
- [71] D. Firesmith, “The need to specify requirements for off-nominal behavior,” *Software Engineering Institute Blog*, Jan 2012.
- [72] Merriam-Webster, “Nominal,” in *Merriam-Webster.com dictionary*, 2023.
- [73] S. ISO, “Robots and robotic devices—collaborative robots (iso-15066: 2016),” *International Organization for Standardization*, 2016.
- [74] S. R. Buss and J.-S. Kim, “Selectively damped least squares for inverse kinematics,” *Journal of Graphics tools*, vol. 10, no. 3, pp. 37–49, 2005.
- [75] S. Scherzinger, A. Roennau, and R. Dillmann, “Forward dynamics compliance control (fdcc): A new approach to cartesian compliance for robotic manipulators,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 4568–4575.
- [76] S. Starke, N. Hendrich, and J. Zhang, “Memetic evolution for generic full-body inverse kinematics in robotics and animation,” *IEEE Transactions on Evolutionary*

- Computation*, vol. 23, no. 3, pp. 406–420, 2018.
- [77] 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.
- [78] R. Luna, I. A. Şucan, M. Moll, and L. E. Kavraki, “Any-time solution optimization for sampling-based motion planning,” in *2013 IEEE international conference on robotics and automation*. IEEE, 2013, pp. 5068–5074.
- [79] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, “Octomap: An efficient probabilistic 3d mapping framework based on octrees,” *Autonomous robots*, vol. 34, pp. 189–206, 2013.
- [80] M. Colledanchise and P. Ögren, *Behavior trees in robotics and AI: An introduction*. CRC Press, 2018.
- [81] R. Ghzouli, T. Berger, E. B. Johnsen, S. Dragule, and A. Wasowski, “Behavior trees in action: a study of robotics applications,” in *Proceedings of the 13th ACM SIGPLAN International Conference on Software Language Engineering*, 2020, pp. 196–209.
- [82] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, “Segment anything,” *arXiv preprint arXiv:2304.02643*, 2023.
- [83] T. Lattimore and C. Szepesvári, *Bandit algorithms*. Cambridge University Press, 2020.
- [84] P. Viola and M. Jones, “Rapid object detection using a boosted cascade of simple features,” in *Proceedings of the 2001 IEEE computer society conference on computer vision and pattern recognition. CVPR 2001*, vol. 1. Ieee, 2001, pp. I–I.
- [85] J. O. Wobbrock, S. K. Kane, K. Z. Gajos, S. Harada, and J. Froehlich, “Ability-based design: Concept, principles and examples,” *ACM Transactions on Accessible Computing (TACCESS)*, vol. 3, no. 3, pp. 1–27, 2011.
- [86] G. Canal, C. Torras, and G. Alenyà, “Are preferences useful for better assistance? a physically assistive robotics user study,” *ACM Transactions on Human-Robot Interaction (THRI)*, vol. 10, no. 4, pp. 1–19, 2021.
- [87] M. Resnick and E. Rosenbaum, “Designing for tinkability,” *Design, make, play: Growing the next generation of STEM innovators*, pp. 163–181, 2013.
- [88] Y. Zhang, Z. Li, H. Guo, L. Wang, Q. Chen, W. Jiang, M. Fan, G. Zhou, and J. Gong, ““ i am the follower, also the boss”: Exploring different levels of autonomy and machine forms of guiding robots for the visually impaired,” in *CHI Conference on Human Factors in Computing Systems*. ACM, 2023, pp. 1–22.
- [89] V. Ranganeni, M. Sinclair, E. Ofek, A. Miller, J. Campbell, A. Kolobov, and E. Cutrell, “Exploring levels of control for a navigation assistant for blind travelers,” in *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, 2023, pp. 4–12.
- [90] J. M. Beer, A. D. Fisk, and W. A. Rogers, “Toward a framework for levels of robot autonomy in human-robot interaction,” *Journal of human-robot interaction*, vol. 3, no. 2, p. 74, 2014.
- [91] 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.
- [92] J. R. Lewis, “The system usability scale: past, present, and future,” *International Journal of Human-Computer Interaction*, vol. 34, no. 7, pp. 577–590, 2018.
- [93] T. A. Spiegel, T. A. Wadden, and G. D. Foster, “Objective measurement of eating rate during behavioral treatment of obesity,” *Behavior Therapy*, vol. 22, no. 1, pp. 61–67, 1991.
- [94] Y. Hamada, A. Miyaji, Y. Hayashi, N. Matsumoto, M. NISHIWAKI, and N. Hayashi, “Objective and subjective eating speeds are related to body composition and shape in female college students,” *Journal of nutritional science and vitaminology*, vol. 63, no. 3, pp. 174–179, 2017.
- [95] M. Hertzum, “Reference values and subscale patterns for the task load index (tlx): a meta-analytic review,” *Ergonomics*, vol. 64, no. 7, pp. 869–878, 2021.
- [96] A. K. Dey, G. D. Abowd, and D. Salber, “A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications,” *Human-Computer Interaction*, vol. 16, no. 2-4, pp. 97–166, 2001.
- [97] “Coding Section GG Self-Care & Mobility Activities Included on the Post-Acute Care Item Sets: Key Questions to Consider When Coding,” *U.S. Center for Medicare and Medicaid Services*.
- [98] M. Rubagotti, I. Tusseyeva, S. Baltabayeva, D. Summers, and A. Sandygulova, “Perceived safety in physical human–robot interaction—a survey,” *Robotics and Autonomous Systems*, vol. 151, p. 104047, 2022.
- [99] A. De Santis, B. Siciliano, A. De Luca, and A. Bicchi, “An atlas of physical human–robot interaction,” *Mechanism and Machine Theory*, vol. 43, no. 3, pp. 253–270, 2008.
- [100] A. Nanavati, E. Gordon, T. A. K. Faulkner, Y. Song, J. Ko, T. Schrenk, V. Nguyen, B. H. Zhu, H. Bolotski, A. Kashyap, and *et al.*, “Lessons learned from designing and evaluating a robot-assisted feeding system for out-of-lab use (supplementary materials),” Dec 2024. [Online]. Available: <https://doi.org/10.17605/OSF.IO/PRSXT>
- [101] I. A. Sucan, M. Moll, and L. E. Kavraki, “The open motion planning library,” *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [102] S. Karaman and E. Frazzoli, “Incremental sampling-based algorithms for optimal motion planning,” *Robotics Science and Systems VI*, vol. 104, no. 2, pp. 267–274, 2010.
- [103] S. Sofaer, “Qualitative methods: what are they and why use them?” *Health services research*, vol. 34, no. 5 Pt 2, p. 1101, 1999.




Fig. 7. A close-up view of the system’s end-effector.

APPENDIX A SYSTEM

This section contains additional details of the system, beyond the core details presented in Sec. III.

A. System Hardware

Fig. 7 shows a close-up of the elements mounted to the robot’s wrist and end-effector: the eye-in-hand system, a fork assembly with a force-torque sensor, and a wireless force-torque transmitter.

B. Co-Design Sessions with the Community Researcher

Many components of the app were co-designed with community researchers. This includes the following.

We workshopped the app state-machine with a community researcher, discussing what the robot and user would do at each step. A key insight from this was the value of including “auto-continue” options so interested users can reduce the number of steps they have to take; this led to the “auto-continue” button before the robot moves to the user’s mouth.

We co-designed the pause, back, resume, and retry options by having the community researcher verbally describe how he would like to control those aspects of robot motion, and creating a mock-up of that in real-time using a picture-editing application. He then gave us feedback, we discussed the pros and cons of the design, and continued until we converged on a user interface.

In a similar fashion, we co-designed the types of transparency the robot should provide the user as it is or is not moving. This involved teaching the community researcher the high-level concepts of robot planning versus motion. This co-design process led us to converge to the robot displaying elapsed time while it is planning but not yet moving, displaying the percent of motion it has completed while it is moving, and displaying a “lock” icon if it is on a screen where it will not move unless the user presses a button.

Finally, we ran a pilot study with a community researcher to investigate how to design the bite selection interface. We first

introduced the community researcher to the concepts of object detection, segmentation, and classification. We then had him load a URL on his phone with 3 mock-ups of different bite selection interfaces: (a) the one described in Sec. III-C1; (b) an interface where the food segmentation algorithm segments several possible bites without a seed point, and the app renders them all to the user as buttons; and (c) an interface where the app renders semantic labels of the food type as buttons. This revealed several pros and cons about each. The first provides the most user control to select a bite, but some assistive technologies make it difficult to select an arbitrary point on an image. The second is more accessible since it has buttons, but can easily become cluttered. The third can be good for users who don’t want to be so involved that they are selecting individual bites, but it provides users little recourse if the robot regularly misdetects a food. Through the discussion, we all converged to starting with the first option, and eventually adding the third option as a choice for users who want less control over their feeding process.

C. Planning

We use MoveIt2 with the Open Motion Planning Library (OMPL) [101] for path planning. This section contains additional details on our use of planning algorithms.

Across both studies, we use the default path length optimization objective, which minimizes the trajectory length in configuration space. Every path planning request has a corresponding time budget: 0.5 seconds for all motions but bite acquisition, which has 2.0 seconds. We generate 5 plans in parallel and hybridize between them [78]. Then, any time remaining in the budget is spent shortcutting the resulting trajectory [78].

The two studies differed in which planning algorithm we used. For Study 1, we use RRT* [102] as the planning algorithm, with all default parameters except `range`, which we set to 3.0 after some informal tuning that sought to balance between planning time and path length within a fixed time budget. For Study 2, in an attempt to speed up the planning times, we switched to RRT-Connect [77] with all default parameters (OMPL by default sets `range` to 20% of the maximum extent of the state space).

Most motions are kinematic plans that use the above pipeline. Of those motions, the ones to hard-coded configurations have goals specified as 6DoF joint goals. For bite acquisition’s motion above the food, the goal is specified as a pose goal for the fork tip, and MoveIt2 samples multiple 6DoF joint goals from the inverse kinematics solver.

A few motions do not involve kinematic plans. For bite acquisition’s motion into the food (approach), we use MoveIt2’s default cartesian planner, which interpolates between the fork tip pose at the start and end, divides it into small intervals, and uses inverse kinematics to get joint configurations for each of those poses. For bite acquisition’s grasp and extract, we use cartesian control to directly execute twists (linear and angular velocities) on the fork tip, using a selectively damped pseudo-inverse Jacobian [74, 75]. For the motions between the “staging” configuration and the user’s mouth, we also use the

aforementioned cartesian control approach to move the fork in a straight line to their mouth.

To avoid collisions, the robot’s static planning scene includes meshes for the user’s head, their body, the furniture they are sitting in (e.g., wheelchair or bed), and the table the food is on. The “body” mesh is a large hull intended to cover diverse body types. The “head” mesh moves and the “body” mesh scales based on the results of face detection. To prevent the robot from generating unnecessarily large motions, tight workspace walls are computed and statically placed in the planning scene to contain the user, the robot (in all of the hard-coded configurations), and some or all of their furniture. To account for un-modeled obstacles, such as user-specific assistive technology (AT) or the user’s laptop, we use the depth image to populate an Octomap [79] at a resolution of 2cm. For plans involving bite acquisition, the robot is allowed to collide with the Octomap and table, because it is intended to come into contact with the food and sometimes the bottom of the plate.

Psychological safety and user comfort is crucial when there is a robot arm moving in close proximity to the user. We promote feelings of user safety with respect to robot motion in a few ways:

- 1) We reject any plan with joint rotations greater than a specified threshold²³, to avoid plans with large swivels that users might find scary or unpredictable.
- 2) For kinematic plans that may move near the face (i.e., the motions to the “staging” and “resting” configurations), we add a wall to the planning scene roughly 0.3m in front of the user’s face. This is intended to keep motions away from the user’s face.
- 3) As mentioned above, the fork moves in a straight line to and from the user’s mouth, to promote interpretability of the robot’s trajectory.

In terms of constraints beyond goal constraints, the system allows for orientation path constraints to be placed on the fork to ensure it remains face-up after acquiring foods. However, since all user-selected foods were skewerable, food falling off due to fork rotations was less of an issue. Thus, we did not use orientation path constraints during either study. Empirically, we found that adding orientation path constraints increased planning time roughly fourfold.

This work optimized planning times to the minimum amount necessary for the system to be deployable out-of-lab. Thus, we do believe the system’s planning can be sped up considerably through, e.g., better tuning of planner parameters, faster inverse kinematics computations, sending multiple goals to the planner, and using MoveIt’s Python bindings as opposed to the ROS2 interface.

D. Food-on-Fork Detection Algorithm

Challenges: Detecting whether there is food on the fork in an RGB image is difficult due to reflections off of the metallic fork and due to the large diversity of food colors.

²³For all motions, the plan was rejected if the sum of joint motion across all joints exceeded 10.0 radians. For acquisition, plans were additionally rejected if joint 1 exceeded $5\pi/6$ or joint 2 exceeded $\pi/2$ radians.

Detecting whether there is food on the fork in a depth image is difficult because the fork is reflective and most of the fork tip is empty space; thus, whether or not the fork is even perceived in the depth image varies depending on lighting condition and objects behind the fork. Both are difficult due to the large variety of food shapes. Finally, detecting food on the fork with the F/T sensor is difficult due to hysteresis errors; after the sensor experiences high force during acquisition, it takes time to regain the level of sensitivity required to detect whether there is food on the fork.

Key Insights: Our approach hinges on two key insights:

- 1) Although food comes in a variety of shapes, *the fork has only one shape*.
- 2) Although the fork may or may not be perceived by the depth camera when there is no food on the fork, *food is always perceived when there is food on the fork*.

Algorithm Overview: Our algorithm uses depth images to memorize the shape of the fork without food. It then learns to predict the likelihood of food on the fork based on the deviation of an input depth image from that memorized shape.

Algorithm Details: Our algorithm operates on de-noised depth images²⁴, converted to pointclouds. During train time, it stores a representative set of points from the “no food on fork” pointclouds—essentially, it memorizes the shape of the fork²⁵. Then, for each pointcloud, it computes the distance between each of its points and the closest point in the stored set, and then takes the 90th percentile of those distances—essentially, this measures how far the farther points in the pointcloud are from the memorized fork shape. Finally, it trains a logistic regression classifier on those distances (x) and whether there is food on the fork (y). During test time, it passes a depth image through the same preprocessing steps, computes the 90th percentile difference between that pointcloud and the stored points, passes that through the logistic regression model, and uses the output as its confidence. If there are < 100 points in the depth image, it outputs nan as its confidence.

Usage: After bite acquisition, if the user has enabled “auto-continue,” the web app toggles on food-on-fork detection and subscribes to its output. If, over the last 3 seconds, the output of food-on-fork is consistently nan or ≤ 0.25 , the web app invokes the action to move the robot above the plate (i.e., acquisition failed). If it is consistently ≥ 0.75 , the web app invokes the action to move the robot to the staging configuration. Else, the web app waits for user input. Similarly, when the robot is at the user’s mouth, if the user has enabled “auto-continue,” the web app toggles on food-on-fork detection and subscribes to its output. If, over the last 3 seconds, the output is consistency ≤ 0.25 , the web app invokes the action to move the robot from the user’s mouth and above the plate (i.e., the user ate the bite from the fork). Else, the web app waits for user input. Note that in the latter case, the web

²⁴During pre-processing, the algorithm crops the depth images to a rectangle around the fork, passes it through a temporal filter that only keeps depth points that are perceived across 5 consecutive images, and passes it through the morphological “opening” operation to remove isolated, noisy points.

²⁵To reduce redundant points, it only stores a point if it is ≥ 1 mm away from all other stored points

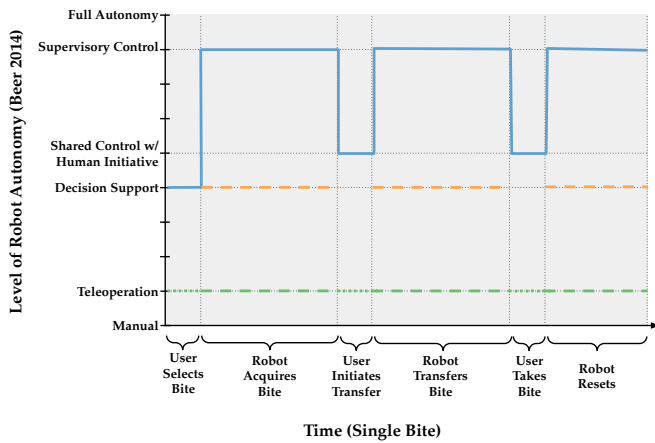


Fig. 9. The levels of control users have access to across each bite.

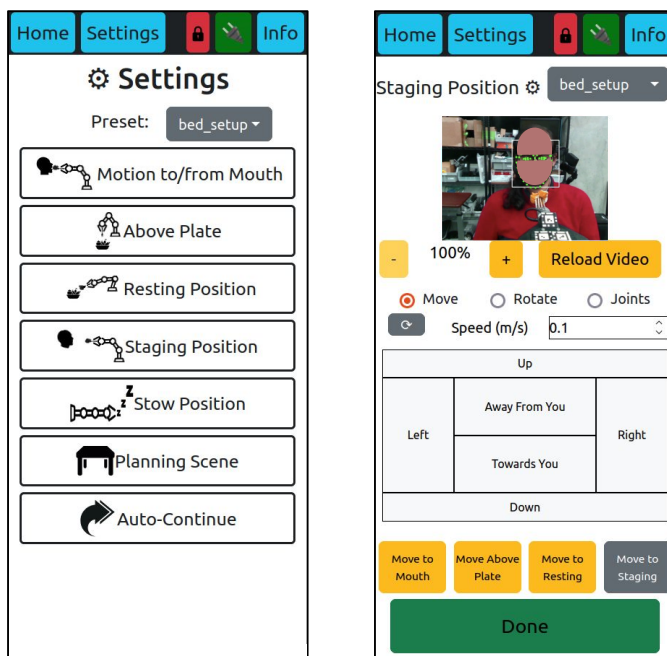


Fig. 8. (Left) The settings menu. (Right) The screen to customize the staging configuration.

app waits for input on nan predictions, because it is possible to get too few points in the pointcloud while the user’s mouth is on the fork.

Limitations: Although this model works, it is susceptible to slight changes in the pose and shape of the fork in the camera frame, which can occur if the fork bends, the camera moves, or the utensil changes. Another approach, which may be more robust, involves using a vision-language model (VLM) to assess whether there is food on the fork [13]²⁶.

E. User Interfaces for Customization

As mentioned in Sec. III-D1, the user is able to customize system parameters through a settings menu in the web app. Fig. 8 Left shows this settings menu. It allows them to customize: properties of the bite transfer motion (distance to

mouth, speeds when approaching the mouth); key robot arm configurations that all motions start or end from; choose which planning scene to use; and change at which points the web app auto-continues. Fig. 8 Right shows the screen the app goes to if the user wants to customize the staging configuration. This illustrates several of the principles from “Designing for Tinkerability” [87] mentioned in Sec. III-D1. For example, the user is given “immediate feedback” by being able to transparently see the impacts of their parameter changes. This happens in two places: (a) seeing how face detection performs in the new configuration, at the top of the screen; and (b) allowing them to invoke actions to/from this configuration to see how the customization impacted robot motion, via the buttons at the bottom of the screen. As another example, the user is given “fluid experimentation” by having access to the

²⁶https://github.com/empriselab/FLAIR/blob/main/bite_acquisition/scripts/food_on_fork.py

full teleoperation interface, seen in the middle of the screen. The user can switch between “move,” “rotate,” and “joints” mode. The former two have 6 buttons, to move the robot in the positive and negative directions of each of the three cartesian motions in that mode. The latter mode has 12 buttons, to move each of the 6 joints in the positive and negative directions.

F. Users’ Multiple Levels of Control

Fig. 9 graphically shows the multiple levels of control users have access to (Sec. III-D2), using the levels of control described in Beer et al. [90]’s framework. The solid line shows the nominal variation of level of control across a single bite. During bite selection, the level of control is “decision support,” since the system presents the user with several masks, which they choose from. During all robot motion, the nominal level of control is “supervisory control,” but the user can drop it down into “decision support” by pausing robot motion, and then into “teleoperation” if so desired. When waiting for the user to initiate bite transfer and to indicate that they are done with the bite, if the user toggles auto-continue on the system is in “shared control with human initiative,” because the robot is automatically deciding whether to continue, but the user can override that decision. If the user has auto-continue toggled off, those stages are at the “teleoperation” level of control, since the user needs to specify when the system should move on. Finally, the user can bypass any of the perception stages altogether by teleoperating the robot from the previous stage onto the next (e.g., fully teleoperating bite acquisition removes the need for bite selection).

APPENDIX B HEALTH & SAFETY PROTOCOLS

Meals can involve health and safety risks. As a result, in both studies the research team strictly adhered to following:

Food Safety: All food was procured from: (a) a restaurant; (b) a grocery store; or (c) homemade by one of the user’s caregivers. The only food preparation the research team did was re-heating, washing, cutting, and/or arranging the aforementioned food, oftentimes with direct input or supervision from the user or their caregiver. All utensils that came in

contact with the food, including the robot's fork were washed with soap and water before every meal. The robot's fork was additionally washed with an alcohol wipe, in front of the user, before they began their meal. The research team washed their hands with soap and water before every meal, and used hand sanitizer before and after touching any food. At any time, the user could request the research team repeat any of the above food safety precautions.

Infection Control: All members of the research team followed the government department of health's COVID-19 prevention and safety guidelines. In addition, masks and hand sanitizers were available during meals, and at any time the user was allowed to ask team members to wear masks and/or take additional health and safety precautions.

Researcher Interventions: Any researcher was allowed to intervene in the meal on: (a) participant request; (b) unexpected or potentially dangerous robot behavior; or (c) perceptible participant distress. They were allowed to take whatever intervention necessary to rapidly resolve the issue, including but not limited to terminating the robot controllers' software or physically powering off the robot.

APPENDIX C COMPARISON TO OTHER ROBOT-ASSISTED FEEDING SYSTEMS

Table I presented a comparison of technical capabilities across contemporary robot feeding systems. This section provides concrete details on the criteria used for each column.

Approximate Cost. For commercial systems, if the system's website or medical tools catalogs mentioned a cost, we used that. If not, we looked through other online sources such as news articles about the technology, crowdfunding campaigns to raise money for the technology, etc. to determine the cost. For research systems, if the paper presented a cost we used that. Otherwise, we followed the aforementioned criteria for all commercially sold components of the system, and added the costs together.

Mounting. This refers to where the robot arm is placed before feeding. We determined this through textual descriptions and pictures of the system.

Autonomous Motion. This refers to whether the robot arm moves autonomously (✓) or whether users have to teleoperate its motions (✗), in nominal scenarios.

(General) Food Detection. This refers to whether the system can detect some (for "Food Detection") or any (for "General Food Detection") bite-sized food items placed in front of the user, without requiring researcher intervention before or during the meal. Note that this column focuses on perception models that can detect masks or bounding boxes around food items, irrespective of whether the perception modules also add semantic labels to those masks. Different robot-assisted feeding works focus on different features for their food detection subsystem. For example, Jenamani et al. [13]'s food detection subsystem does not provide general food detection, because researchers must seed it with a list of food items on the plate. However, it does provide additional features not encompassed by this paper's bite selection subsystem.

Specifically, their system can semantically segment foods and can segment non-bite sized food items (e.g., spaghetti or mashed potatoes).

Face Detection. This refers to whether the system can autonomously detect the user's face and mouth.

Collision Detection / Avoidance. "Collision Detection" refers to whether the system can detect a collision once it has occurred (and stop/modify its motion accordingly). "Collision Avoidance" refers to whether the system can preemptively avoid possible collisions, e.g., through its motion planning.

Portable & Self-contained. This refers to whether the system can be moved, with the user and caregiver, to the varied locations that users may eat in: e.g., at home, at a restaurant, at an outdoor picnic table, etc. Reasons a system may not be portable & self-contained include: the system requires wall power; the system is too heavy to move; the system has too large of a footprint to exist in the diverse environments people eat in; the system has wires that stretch across robot joints, which can get tangled, restrict robot motion, and be trip hazards.

User Can Stop / Restart Robot Motion. The former refers to the user's ability to stop the robot at any time, for any reason. The latter refers to the user's ability to restart the robot once they have stopped it. For example, a system that provides only an emergency stop (e-stop) button that requires a researcher to restart the system would satisfy the former but not the latter criterion. In contrast, a system that allows the user to stop robot motion via, for example, an app, and subsequently restart robot motion, satisfies both criteria.

Customizable Robot Motion. This refers to whether users can customize the autonomous motions the robot takes, for example customizing its speed, how close it gets to their mouth, the path it takes to get to their mouth, etc. Note that some systems that don't provide customizable robot motion do provide other forms of customizability, e.g., Jenamani et al. [13] allow users to customize the sequence the robot provides them bites in.

Multiple UI Modalities. This refers to whether the system allows users to use multiple interface types to interact with it, depending on their impairments and preferred assistive technologies (ATs). A system whose user interface is on a general-purpose computing device such as a smartphone or laptop implicitly satisfies this, since general-purpose computing devices typically are designed to be compatible with diverse ATs.

Note that Table I refers to the technical capabilities, as stated in the papers or websites. Table X, below, presents the *demonstrated* capabilities, i.e., the system's performance as demonstrated in a published user study.

APPENDIX D STUDY 1: MULTI-USER, ON-CAMPUS STUDY

This section contains additional details of Study 1, beyond the core details presented in Sec. IV.

Each user participated in a 30 minute virtual meeting before the study, followed by a 90 minute in-person session, where they ate an entire meal of their choice in an out-of-lab location.



Fig. 10. An example of our in-person meal setup in a public cafeteria.

They received \$25/hour compensation for their time, as well as compensation for travel to and from the study venue.

A. Virtual Session Details

We asked each participant to describe any assistive technologies that they use:

- 1) Do you regularly use a smartphone or tablet?
- 2) If so, how do you interact with the smartphone or tablet?
- 3) If you use an assistive technology (AT) to interact with the smartphone or tablet, where is that AT mounted?
- 4) If you use an AT to interact with the smartphone or tablet, how do you interact with the AT?
- 5) For the in-person study, would you be able to bring your smartphone or tablet, along with your preferred AT to interact with it?
- 6) Do you have other ATs mounted around your head or chair?

We also asked participants about their food preferences:

- 1) Do you have any allergies or other dietary restrictions we should be aware of?
- 2) As the in-person portion of the study will involve eating a full meal, what would you like to eat?

If the participant was having trouble identifying desired food items, we provided sample food items from the following list:

- **Proteins:** sandwich meats, chicken tenders, cheeses, blocks of baked or fried tofu, etc.
- **Vegetables / Fruits:** salads, roasted vegetables, crudites, fruit salad, etc.
- **Starches:** Potatoes, Rice, Bread, Noodles, etc.

Finally, we asked them about transportation logistics:

- 1) How do you anticipate coming to [the study venue]? What expenses are associated with your transportation?
- 2) Will someone be coming with you?
- 3) Is there anything else we can do to make the in-person study accessible to you?

B. In-Person Session Details

Fig. 10 shows a representative in-person study setup, for the meal in the cafeteria with CR2. Every in-person session had one camera zoomed in on the participant's phone or tablet, and another that captured the participant, robot, and social dining partner.

1) *System Introduction:* A researcher introduced each participant to the system as follows:

- 1) Set the system on a tripod next to the participant's wheelchair.
- 2) Mounted the emergency stop button in a location the participant could reach, and explained how to use it to stop the robot in the case of an emergency.
- 3) Assisted the participant in connecting their phone/tablet to the system's WiFi network and opening the system web app in a browser tab.
- 4) If needed, assisted the participant in setting up any assistive device to interact with the web application.
- 5) Walked the participant through completing one bite using the web application. Demonstrated safety features, such as the F/T sensor's ability to stop the robot when unexpected forces occur.
- 6) If requested by the participant, walked them through customizing how close to their mouth the robot gets.
- 7) Performed any necessary system adjustments, such as moving the tripod or plate, throughout the above process.

2) *Pre-Post-Meal Questions:* After the practice bite but *before* the full meal started, we asked the participant the following five-point Likert scale questions, where [AID_TYPE] is replaced by either "my caregiver," "my self," or both, depending on how the participant eats on a regular basis. At the end of the meal, we asked them these questions again, with [AID_TYPE] replaced with "the robot" and question tense shifted to past. Each answer was on the scale: Strongly Disagree, Disagree, Neutral, Agree, Strongly Agree.

- 1) When I eat with [AID_TYPE], I get my next bite when I want it, without waiting or feeling rushed.
- 2) When I eat with [AID_TYPE], I decide what food I want in my next bite.
- 3) When I eat with [AID_TYPE], I feel a sense of independence.
- 4) When I eat with [AID_TYPE], the meal requires a lot of mental energy.
- 5) When I eat with [AID_TYPE], the meal requires a lot of physical energy.
- 6) When I eat with [AID_TYPE], I am confident that I will remain safe during the entire meal.
- 7) When I eat with [AID_TYPE], I am confident that I will remain clean during the entire meal.

We then asked a final question (both pre- and post-meal): "How comfortable are you with the idea of being fed by a robot?" on the scale: Very Uncomfortable, Uncomfortable, Neutral, Comfortable, Very Comfortable.

Note that to accommodate their impairments, we asked all quantitative questions to participants verbally. Before asking any questions after the meal, we reminded participants that negative feedback is also very helpful for us to know how to improve the system.

3) *Cognitive Workload (NASA-TLX):* Because we asked all quantitative questions to participants verbally, we modified the wording of the NASA-TLX [91] to be:

- 1) On a scale of 0-20, how mentally demanding has the task been? 0 is very low, 20 is very high.




Fig. 11. Images of each user's plate of food for Study 1, taken at various points in the meal. Several users requested we serve them more food.

- 2) On a scale of 0-20, how physically demanding has the task been? 0 is very low, 20 is very high.
- 3) On a scale of 0-20, how hurried or rushed has the pace of the task been? 0 is very low, 20 is very high.
- 4) On a scale of 0-20, how successful have you been at accomplishing the task? 0 is failure, 20 is perfect²⁷.
- 5) On a scale of 0-20, how hard have you had to work to accomplish your level of performance? 0 is very low, 20 is very high.
- 6) On a scale of 0-20, how insecure, discouraged, irritated, stressed, or annoyed have you been? 0 is very low, 20 is very high.

4) *System Usability Scale (SUS)*: We asked the exact questions of the system usability scale [92], on a 5-point Likert Scale: Strongly Disagree, Disagree, Neutral, Agree, Strongly Agree.

- 1) I think that I would like to use this system frequently.
- 2) I found the system unnecessarily complex.
- 3) I thought the system was easy to use.
- 4) I think that I would need the support of a technical person to be able to use this system.
- 5) I found the various functions in this system were well integrated.
- 6) I thought there was too much inconsistency in this system.
- 7) I would imagine that most people would learn to use this system very quickly.
- 8) I found the system very cumbersome to use.
- 9) I felt very confident using the system.
- 10) I needed to learn a lot of things before I could get going with this system.

C. Between-Study System Patches

The order of participants in the study was: P1, P2, P3, CR2, P4, and P5. There were 6 day gaps each between P1 and P2 and between CR2 and P4²⁸. This gave us the time to

²⁷When reporting this value (Table VII), we flip it (i.e., 0 is perfect) to align with the original NASA-TLX.

²⁸There was a 0 or 1 day gap between all other participants

conduct system patches to address bugs revealed in previous meals. The applied system patches were:

- 1) Between P1 and P2, we:
 - a) Sped up all joint velocity limits by 66% on all kinematic motions, and sped up the cartesian motion velocity limit when moving to/from the mouth by 20%, from 0.1 to 0.12m/s.
 - b) Addressed a bug where food detection's depth readings would get skewed if the fork partially overlapped the bite.
 - c) Addressed a bug where if the user moves the arm from their mouth back to the staging configuration while "auto-continue" is checked, it will subsequently move right back to their mouth.
 - d) Modified the "MoveToMouth" action to reset the Octomap before starting; this is to account for phantom obstacles that accrue over time.
 - e) Added the ability to zoom into the robot's camera feed during bite selection.
 - f) Relaxed the deviation from goal position that is accepted when the robot is moving to the user's mouth from 0.5cm to 2.5cm.
 - g) Addressed a bug where sometimes the joint state publisher's message timestamps are before the camera's, leading to failures in transforming between those frames of reference.
- 2) Between CR2 and P4, we:
 - a) Added an auto-restart process manager around the WebRTC signalling server, to address a known bug in a dependency that sometimes causes a segmentation fault.
 - b) Addressed a bug where food detection would return masks that have few valid depth readings (e.g., due to being too close to the image edge).
 - c) Addressed a bug where sometimes the arm would make a large, unnecessary swivel to get from one configuration to another, by rejecting plans where joints rotated being a certain threshold

- d) Addressed a hardware issue where the bolts in one finger had loosened, resulting in the gripper holding onto the fork asymmetrically.
- e) Replaced the e-stop button’s adapter due to regular wear and tear.
- f) Added a recovery behavior where the robot arm raises itself up 1cm if motions fail during bite acquisition (to prevent the case where the fork is left in contact with the table, causing all future actions to fail due to an unexpectedly high force sensor reading).
- g) During bite acquisition, had the robot plan both the motion above the plate and into the food before executing. That way, any planning failures will happen while the robot is still above the plate, as opposed to after it starts moving.

Thus, P1 experienced the system with neither System Patch, P2, P3, CR2 experienced the system with System Patch 1, and P4 and P5 experienced the best version of the system, with both System Patches. Importantly, the system P1 experienced had the robot arm moving up to 66% slower than the system all other participants experienced.

D. Data Analysis

For the objective data analysis, one researcher watched the videos recorded from every sessions and tagged the timestamps of all key events in the video. A key event was defined as when the user interacted with the web app, the robot started/stopped moving, an off-nominal scenario occurred that was resolveable without researcher intervention, and an off-nominal scenario occurred that required researcher intervention. In addition, every time a bite acquisition or motion to the user’s mouth ended, the researcher tagged the food type and whether or not it was successful. This resulted in a complete time profile of the meal, as well as a complete log of bite acquisition and transfer success rates. All the annotated data, with key events and timestamps per participant over the entire meal, can be found in Supplementary Materials, along with a codebook describing each key event.

For the subjective data analysis, we scaled all 5-point Likert scales to integers in the range $[-2, 2]$. For the NASA-TLX, we followed Hertzum [95]’s procedure of scaling each subscale to $[0, 100]$ and averaging them. For the SUS, we followed Lewis [92]’s procedure of setting missing values to “neutral,” flipping negative subscales, transforming every subscale into the range $[0, 10]$, and summing them.

For both objective and subjective data, due to the small sample size, we do not analyze for statistical significance.

E. What Parts of the System Were(n’t) Evaluated

Most of our system was evaluated in Study 1: the robot, fork holder and F/T sensor, e-stop button, all the robot code, and the web app. However, a few components were not included in Study 1, and later evaluated in Study 2. First, the only aspect of customizability included in Study 1 was: (a) customizing how close to the user’s mouth the robot stops; and (b) toggling auto-continue on/off during face detection. Notably, since the other two auto-continues were not included in this evaluation, neither was the food-on-fork detection module. Second, we did

not give users access to a teleoperation interface to control the robot. Thus, during robot motion, they had the “supervisory control” to drop the robot into “decision support,” but could not drop it into the teleoperation level of control (Sec. III-D2). Finally, in order to improve reliability in the study, we paused online learning for deciding which action the robot should take (Sec. III-C2). As a result, the robot only executed one fixed motion primitive throughout the entire meal (except with P3, who requested we change the primitive so the robot could better acquire strawberries).

F. Results

1) *Plates of Food*: All food was acquired from local restaurants and grocery stores, based on the foods they had requested in the virtual session. We put these foods on a variety of plates, as shown in Fig. 11. Notably, the plates were of various colors and patterns (red, white, blue, patterned), shapes (flat, deep), and materials (paper, ceramic, melamine). The plates also did not necessarily have good color contrast with the food on it, as evidenced by the pink salmon on the red plate. This demonstrates the generalizability of the bite selection approach to diverse plate types.

2) *Bite Acquisition*: Table V shows the complete bite acquisition data, disaggregated by food type. This reveals that for diverse food types, from sandwich bites to broccoli to pizza to salmon, the system is close to users’ threshold acquisition success rate of 80%²⁹. Further, note that because the online learning system was paused during this evaluation (Appendix D-E), the robot was unable to learn from its failures and try different actions.

3) *Time Profile Comparison Across Participants*: Fig. 12 shows a comparison of the time each participant spent in each stage of robot-assisted feeding.

Although one might think that the reason P1 took the longest time per bite was that the robot was up to 66% slower for him than for other participants (Appendix D-C), Fig. 12 reveals that the difference actually has to do with assistive technology. Consider when the robot was “Waiting for User (Bite Done).” P1 used voice control to interact with his phone, which meant that he: (a) could not interact with his phone while chewing; and (b) could not interact with his phone when he or others were talking. As a result, after the robot delivered a bite to his mouth, he waited until he was done chewing (which a spot check revealed took around 30 seconds per bite) and until there was a pause in the conversation before sending the robot back.

The other person who used mouth-based assistive technology, CR2, used a mouth joystick to interact with his phone. This meant that he could not interact with his phone while chewing³⁰, nor while he was talking, but could do so while

²⁹Although Bhattacharjee et al. [19] mention a 70% threshold bite acquisition success rate in the paper, we re-analyzed the raw data from that work, which was shared with us by the authors. The number presented in the paper is the arithmetic mean; however geometric mean tends to be more representative when the numbers are proportions or ratios. The geometric mean of the user data is 80%. Thus, we use 80% as the threshold bite acquisition success rate, which also aligns with Gordon et al. [9].

³⁰CR2’s mouth joystick requires him to suck air out of a straw to “click” a button, which can be a choking hazard if done while chewing

Acquisition Success Rate Per Food

P1	Pizza: 0.78 ($14/18$); Broccoli: 1.0 ($1/1$)
P2	Chicken Tenders: 0.85 ($11/13$); Broccoli: 0.75 ($9/12$); Pasta: 0.33 ($4/12$)
P3	Sandwich: 0.94 ($16/17$); Strawberry: 0.24 ($4/17$); Brownie: 1.0 ($7/7$); Grape: 1.0 ($2/2$); Pineapple: 1.0 ($1/1$); Cantaloupe: 1.0 ($1/1$)
P4	Grilled Chicken: 1.0 ($13/13$); Potato: 1.0 ($12/12$); Cauliflower: 0.56 ($5/9$)
P5	Mac and Cheese: 0.7 ($7/10$); Brussel Sprouts: 0.86 ($9/7$); Salmon: 0.71 ($5/7$); Mushroom: 1.0 ($2/2$); Donut: 1.0 ($2/2$); Chocolate Cake: 1.0 ($1/1$)
CR2	Strawberry: 0.75 ($3/4$); Tofu: 1.0 ($3/3$); Grape: 1.0 ($3/3$); Cantaloupe: 1.0 ($2/2$); Broccoli: 0.5 ($1/2$); Bagel: 0.0 ($0/2$); Pineapple: 1.0 ($1/1$); Beef: 1.0 ($1/1$)

TABLE V

BITE ACQUISITION SUCCESS RATES DISAGGREGATED BY FOOD TYPE. MOST SUCCESSFUL FOODS (≥ 3 BITES) ARE HIGHLIGHTED.




Fig. 12. A box-and-whisker plot showing the 25th, 50th, and 75th percentiles (vertical lines in box) for each of the 6 stages of robot-assisted feeding, for all users' successful bites. Notches (diagonal lines) show the 95% confidence interval around the median. Outliers excluded.

others were talking. This added a layer of parallelization to the meal, which enabled CR2 to send the robot away from his mouth faster than P1 could.

The remaining users all interacted with their phone by using a stylus or touch, so they could interact with it while chewing or conversing. This added an additional layer of parallelization to the process, as those users could send the robot away from their mouth even as they chewed.

The impact of assistive technology can also be seen during "Bite Selection." Voice control is designed to click buttons. Although it is possible to tap arbitrary points, that involves a time-consuming process of zooming into a multi-layered

grid to select the desired point to tap. As a result, it took P1 much longer than other participants to select his desired bite (median: 22 sec). CR2's mouth joystick is a pointer-based interface, but involves moving a cursor along the screen, which takes more time than directly tapping a point on the screen. Thus, CR2 took the second-most time on bite selection (median: 10 sec), followed by the remaining participants who used touch to interact with their devices (median: ≤ 6 sec).

This reveals the importance of not only ensuring the system works for diverse assistive technologies, but also considering how those assistive technologies impact user experience.



Fig. 14. (a) P3’s custom 3D-printed self-feeding tool for grasping onto and maneuvering a fork. (b) P5’s self-feeding tool for strapping a fork to her hand.

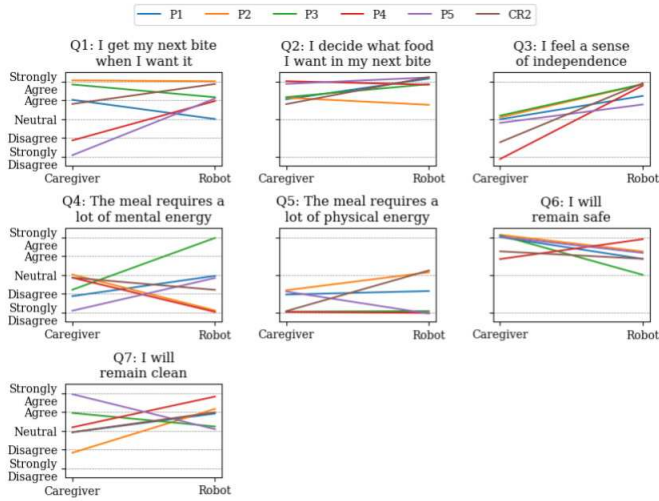


Fig. 13. All questions in users’ self-reported comparison of eating with caregivers vs. the robot.

4) *Time Profile Comparison to Caregiver Feeding*: To compare the time profile of robot-assisted feeding to caregiver feeding, we analyzed a video of CR1 being fed a lunch of mixed berries and a protein bar by their caregiver. The motion to acquire a bite and transfer it to the user occurred in one smooth swoop, taking 1 – 3 seconds. While the care recipient chewed, the caregiver acquired the next bite and was ready as soon as the community researcher finished chewing (around 15 seconds). This reveals a large space to improve our system’s bite duration, both in terms of speed and parallelism. However, we note that not all people with motor impairments want their robot to feed them as fast as their caregivers: some feel that their caregiver’s speed puts pressure on them [6]. The timestamped annotations from the video of CR1’s meal can be found in Supplementary Materials.

5) *Robot-assisted vs. Caregiver Feeding*: Table VI and Fig. 13 shows participant responses to all of the pre-post questions. Highlighted values refer to the aid type (caregiver, robot, or self) that performed highest for that participant on that question. As can be seen, the robot consistently performed as well or better than caregivers for “I feel a sense of independence” and “I decide what food I want in my next bite,” and mostly outperformed caregivers on “I get my next bite when I want it.” Further, note that for P4 and CR2, the robot outperforms the caregiver on nearly every question.

6) *Cognitive Workload (NASA-TLX)*: Table VII show participant responses to all NASA-TLX subscales, as well as their overall cognitive workload. We compare our results

to Hertzum [95]’s baseline mean and standard deviation, computed from 41 studies that used the NASA-TLX to evaluate cognitive workload during studies with “special-needs users.” This reveals that for all participants but P3, our system involved less cognitive workload than the average study with “special-needs users” from Hertzum’s sample.

Note that the amount of physical demand required depended on the assistive device the participant used to interact with their phone or tablet. For example, in order to use her stylus to reach the top part of her phone (for bite selection), P2 needed to leverage the left armrest of her wheelchair to pull herself to the left, thereby getting the right angle to click with the stylus. Similarly, for CR2 to use his mouth joystick to tell the robot to move away from his face after he ate the bite, he had to move his head around the fork to the mouth joystick, which could be a complicated maneuver. On the other hand, P1’s voice control required no additional physical demand for him to interact with his phone.

7) *System Usability Scale (SUS)*: Table VIII shows participants’ ratings for the subscales of the System Usability score (SUS) and the final score. We compute usability grades using the curved grading scale developed from a meta-analysis of over 400 studies that used the SUS [92]. As can be seen, 3 out of the 5 participants—P2, P4, and P5—and CR2 gave the system usability’s ratings that corresponded to average or above-average usability (C or above).

8) *Ranking Aspects of Feeding Systems*: Table IX show user’s responses when asked to rank the top three aspects of robot-assisted feeding systems to improve. 4/5 participants (excluding CR2) put speed as their top choice. Other aspects that commonly occurred across users were portability and independent use, safety, and customizability. These rankings provide pointers towards the most pressing areas of improvement necessary for the robot-assisted feeding research community.

9) *Self-Feeding Tools*: Two participants had custom tools to enable them to feed themselves (Fig. 14). P3 had a custom 3D printed fork holder that enabled him to maneuver the fork to the plate and to his mouth, by using the table as a fulcrum, all while keeping his hands close to his lap and his face above the plate. P5 had a custom-designed strap that attached a fork tip to her hand, so she could use her arm to move the fork tip into foods and move them to her face. Importantly, all the foods P3 ate in the study were also skewerable by his self-feeding device. For P5, some pieces of brussel sprout were skewerable by her self-feeding device—if not too hard—and some pieces of fish were skewerable—if not too flaky—but overall the robot did feed her bites that she felt she may not have been able to skewer with her self-feeding device.

G. Comparison to Other Robot-Assisted Feeding Systems

³¹One user fed himself 130 bites over 6 sessions (Avg: 21.7), while the other 8 fed themselves 20 bites over 1 session.

³²This bimanual manipulator holds the bowl in one hand and the utensil in the other, reducing the distance to traverse for acquisition and transfer.

³³For these works, the one out-of-lab environment was an in-home environment, like that in our Study 2.

	P1		P2		P3			P4		P5			CR2	
	Care-giver	Robot	Care-giver	Robot	Care-giver	Robot	Self	Care-giver	Robot	Care-giver	Robot	Self	Care-giver	Robot
I get my next bite when I want it (↑)	1	0	2	2	2	1	2	-1	1	-2	1	2	1	2
I decide what food I want in my next bite (↑)	1	2	1	1	1	2	2	2	2	2	2	2	1	2
I feel a sense of independence (↑)	0	1	0	2	0	2	2	-2	2	0	1	1	-1	2
The meal requires a lot of mental energy (↓)	-1	0	0	-2	-1	2	0	0	-2	-2	0	-2	0	-1
The meal requires a lot of physical energy (↓)	-1	-1	-1	0	-2	-2	1	-2	-2	-1	-2	-2	-2	0
I am confident that I will remain safe (↑)	2	1	2	1	2	0	2	1	2	2	1	2	1	1
I am confident that I will remain clean (↑)	0	1	-1	1	1	0	1	0	2	2	0	0	0	1

TABLE VI

USER RESPONSES TO THE PRE-POST QUESTIONS. HIGHLIGHTED VALUES ARE THE BEST PER-QUESTION PER-USER.




Fig. 15. The full timeline of P1's meal.

	P1	P2	P3	P4	P5	CR2
I think that I would like to use this system frequently. (↑)	1	1	-2	2	-1	2
I found the system unnecessarily complex. (↓)	1	0	0	-2	-2	-1
I thought the system was easy to use. (↑)	0	1	1	2	1	2
I think that I would need the support of a technical person to be able to use this system. (↓)	-1	-1	2	-1	-2	1
I found the various functions in this system were well integrated. (↑)	0	1	0	1	2	1
I thought there was too much inconsistency in this system. (↓)		0	1	-2	-2	-1
I would imagine that most people would learn to use this system very quickly. (↑)	1	1	-1	2	2	2
I found the system very cumbersome to use. (↓)	-1	-1	-1	-2	1	-1
I felt very confident using the system. (↑)	1	1	0	2	1	2
I needed to learn a lot of things before I could get going with this system. (↓)	-1	1	-1	-1	-1	-2
System Usability Score (SUS) (↑)	62.5	65	42.5	92.5	77.5	82.5
SUS Grade (↑)	D	C	F	A+	B+	A

TABLE VIII

PARTICIPANTS' USABILITY RATINGS FOR THE ROBOT-ASSISTED FEEDING SYSTEM. HIGHLIGHTED RATINGS ARE AT-OR-ABOVE AVERAGE.

	First	Second	Third
P1	Speed	Robustness to errors	Easy user interface
P2	Speed	Robustness to errors	Safety
P3	Portability & independent operation	Customizability	Non-intrusive in daily routine
P4	Speed	Portability	Customizability
P5	Speed	Speed	Customizability
CR2	Safety	Portability & usability	Robustness to errors

TABLE IX
USERS' RANKINGS FOR THE MOST IMPORTANT ASPECTS OF ROBOT-ASSISTED FEEDING SYSTEMS TO WORK ON.

	P1	P2	P3	P4	P5	CR2	Baseline
Mental Demand (↓)	25	25	35	25	25	15	43 ± 16
Physical Demand (↓)	0	50	20	25	15	25	27 ± 12
Temporal Demand (↓)	0	25	10	50	0	25	33 ± 17
Performance (↓)	30	25	40	0	25	10	44 ± 21
Effort (↓)	25	50	75	10	50	25	42 ± 17
Frustration (↓)	25	0	50	10	0	15	31 ± 15
Cognitive Workload (↓)	17.5	29.2	38.3	20	19.2	19.2	37 ± 11

TABLE VII

PARTICIPANTS' COGNITIVE WORKLOAD (NASA-TLX) AFTER EATING WITH THE ROBOT. BASELINE IS MEAN ± STANDARD DEVIATION [95]. HIGHLIGHTED VALUES ARE LESS THAN BASELINE MEAN.

Robot	Num Food Types	Num End-Users Fed (Out-of-Lab)	Num Fed Bites Per User Per Session	Avg. Bite Duration Per User (sec)	User-Decided Meal?	Entire Meal?	Arbitrary Plate?	Num Out-of-Lab Environments	Validated Metrics?	Avg. TLX Score
Song and Kim [20] / [26]	- / -	7 (-)	-	-	-	✓	✗	-	-	-
Song et al. [26]	-	14 (-)	-	-	-	✓	✗	-	-	-
Park et al. [21]	8	9 (1)	20–21.7 ³¹	41–78 ³²	✗	✗	✗	1 ³³	TLX	18.6
Nguyen [22]	1	1 (1)	10	330	✓	✗	✓	1 ³³	-	-
Bhattacharjee et al. [19]	3	10 (0)	15	90	✗	✗	✗	0	-	-
Jenamani et al. [15]	2	13 (1)	7–13	-	✗	✗	✗	1 ³³	-	-
Jenamani et al. [13]	5	1 (0)	-	-	✗	✓	-	0	-	-
This paper’s Study 1	22	6 (6)	14–31	62–165	✓	✓	✓	3	TLX, SUS	23.9

TABLE X

COMPARISON BETWEEN THE DEMONSTRATED CAPABILITIES FROM STUDY 1 VERSUS OTHER ROBOT-ASSISTED FEEDING SYSTEMS’ DEMONSTRATED CAPABILITIES

Table X compares the demonstrated capabilities of our system in Study 1 to the demonstrated capabilities of other research systems that had an evaluation with people with motor impairments. This shows that our system is the first to feed users entire meals of their choice in multiple out-of-lab environments. Additionally, our system fed users over $2\times$ more food types than others (full list in Table V). However, the upper range of bite duration in our system is slower than most others.

H. Off-Nominal Scenarios Per-Participant

This section contains a description of every off-nominal scenario, researcher intervention, bite acquisition failure, and successful bite whose duration was an outlier ($\geq 1.5 \cdot \text{IQR}$) relative to that participant’s other successful bites.

I. P1 Study Details

Fig. 15 shows the timeline of P1’s meal. Notable events include:

- **Bite 0 (User-resolved Off-nominal):** The robot stopped too far from P1’s mouth (likely due to a phantom obstacle in the Octomap; addressed with System Patch 1d). The user was able to get it to move the rest of the way to their mouth by clicking “retry” on the app.
- **Bite 1 & 6 (User-resolved Off-nominal):** Although the robot moved close enough to the user’s mouth for them to eat the bite, from the robot’s perspective the action failed because it encountered an Octomap collision slightly before its goal. This is because the user leaned forward as the robot was coming in. The user resolved this by clicking “retry,” letting the robot move the remaining short distance it wanted to, and then having it go back. (Addressed with System Patch 1f)
- **Bite 2 (User-resolved Off-nominal):** The robot failed to plan to move above the food (addressed with System Patch 1g); user resolved by going back above the plate and re-selecting the bite.
- **Bite 5 (User-resolved Off-nominal):** As with Bite 1, the robot moved close enough to the user’s mouth, but thought it encountered an error (addressed with System Patch 1f). This time, instead of clicking “retry” the user mistakenly clicked “back,” which took the robot back to the staging configuration. From there, “auto-continue” caused the robot arm to move back to his face (addressed with System Patch 1c), and then the user had it move back above the plate.

- **Bite 7, 8, & 11 (Bite Acquisition Failures):** The robot misperceived the depth of pieces of pizza that were partially obscured by the fork, and therefore moved down too little towards the food. Addressed with System Patch 1b.
- **Bite 11-12 (Researcher Intervention):** With the participant’s consent, researchers moved the plate so that none of the pieces of food were partially obscured by the fork in the “above plate” configuration, to avoid the aforementioned issue with food depth misperception.
- **Bite 18 (User-resolved Off-nominal):** Although bite acquisition visually succeeded, the robot thought it encountered an error. The user overcame this by having the robot move back above the plate, and then clicking “skip acquisition” on the bite selection page of the web app.
- **Bite 19 & 19 (User-resolved Off-nominal):** First, none of the detected masks were the bite P1 wanted, so he selected another point on the robot’s camera feed. Second, P1’s phone mistakenly interpreted a number he was saying as part of the conversation as a button click. Thus, his phone opened the live video view of the web app. When he realized, P1 closed out of it.
- **Bite 19 (Outlier Bite Duration):** The user continued the conversation for around 2 minutes with the robot in front of his face, before realizing and sending it back above the plate.

The additional system patches after P1 were based on his feedback that the robot should be sped up (System Patch 1a) and that the plate in the camera view is too small for bite selection (System Patch 1e).

J. P2 Study Details

Fig. 16 shows the timeline of P2’s meal. Notable events include:

- **Bite 0, 12 (Bite Acquisition Failure):** The user’s selected mask was two bites of food together, leading the fork to pierce between the two.
- **Bite 1 (Researcher Intervention):** The force-torque sensor disconnected from WiFi. To address this, researchers lifted the backpack containing the router up off of the floor, restarted the force-torque sensor’s code, and restarted the physical force-torque sensor.
- **Bite 3-4 (Researcher Intervention):** Researchers restarted the WebRTC signalling server, which crashed due to a segmentation fault (addressed in System Patch 2a).




Fig. 16. The full timeline of P2’s meal.




Fig. 17. The full timeline of P3’s meal.

- **Bite 2-3, 4, 9, 10 (Bite Acquisition Failure), Researcher Intervention:** The fork wasn’t centered on the bite, so researchers nudged the fork in the gripper to get it to better align with the robot’s URDF model (addressed in System Patch 2d).
- **Bite 6, 7, 20, 24, 28 (Bite Acquisition Failure):** The pasta rolled away as the robot was skewering it.
- **Bite 8, 22-23 (Researcher Intervention):** Researchers restarted the force-torque sensor software due to it ceasing to send or receive messages.
- **Bite 19, 27 (Bite Acquisition Failure):** The food segmentation algorithm only detected a subset of the pasta helix, not the whole piece, leading the robot to approach it not perpendicular to the piece.
- **Bite 24 (Researcher Intervention):** Researchers reset the software because the app and robot stopped communicating.
- **Bite 25 (Outlier Bite Duration):** Due to paying attention to the conversation, the participant waited after initiating bite selection and after the robot moved to the resting pose, thereby extending this bite.
- **Bite 28-29 (User-resolved Off-nominal):** Force-torque sensor lost WiFi connection, but regained it shortly without requiring researcher intervention. The user re-initiated bite acquisition after force-torque connection was restored.
- **Bite 32 (Researcher Intervention):** Researchers restarted the force-torque sensor hardware, due to an issue with receiving UDP packets.
- **Bite 32-33 (User-resolved Off-nominal):** The robot’s camera feed did not render on the app in bite selection. The user clicked the “reload video” button and then it did.
- **Bite 3, 13, 23-24, 42 (User-resolved Off-nominal):** None of the masks returned from food segmentation aligned with the user’s desired bite. To address this, they re-invoked food segmentation with a new seed point.
- **Bite 6 (Researcher Intervention):** During acquisition, the fork missed the food and hit the plate. For some reason (perhaps being too close to a singularity) the extraction motion out of the food failed, leaving the robot in contact with the plate. This caused all future actions to fail due to the robot experiencing a higher force than the threshold. To address this, researchers briefly manually lifted the robot arm, getting it out of contact with the table, while the participant invoked the “move above plate” action on the app. Addressed in System Patch 2f.
- **Bite 7, 8, 9, 24 (Bite Acquisition Failure):** Due to the strawberry being extremely soft, it slid off the fork as the robot was lifting the fork up.
- **Bite 10-11, 11, 23-24 (Researcher Intervention):** Researchers nudged the fork in the gripper, so it better aligns with the robot’s URDF model (addressed in System Patch 2d).
- **Bite 12 (User-resolved Off-nominal):** Robot was unable to find a plan to move into the food. User clicked “back” and re-selected their desired food item. Partially addressed in System Patch 2g.
- **Bite 12-13 (Researcher Intervention):** With participant consent, researchers moved the plate to be more centered on the robot arm in its “above plate” configuration, to increase the likelihood of motion success.
- **Bite 15, 34 (Bite Acquisition Failure):** Robot arm pushed the strawberry out of the way as it was descending into it, because the curve of the strawberry aligned with the curve of the fork tines.
- **Bite 16, 33, 45 (Bite Acquisition Failure):** Robot skewered the strawberry, but it fell off as the robot was moving to the “resting” configuration.

K. P3 Study Details

Fig. 17 shows the timeline of P3’s meal. Unlike all other participants, P3 preferred to sit near the front of their wheelchair; this resulted in a much shorter distance to/from his mouth. Notable events in this meal include:




Fig. 18. The full timeline of CR2’s meal.

- **Bite 20 (Researcher Intervention):** Robot did an extremely large, multi-part motion when moving above the plate. A researcher attempted to click the emergency stop button, which didn’t register, so another researcher terminated the controllers. Swivel issue addressed in System Patch 2c, e-stop button issue addressed in System Patch 2e. Researchers restarted the software afterwards and the participant continued.
- **Bite 23 (Bite Acquisition Failure):** Robot arm was off-center and missed the strawberry (addressed in System Patch 2d).
- **Bite 33 (User-resolved Off-nominal):** User mistakenly initiated a bite transfer after a failed bite acquisition. He promptly paused and had the robot arm go back.
- **Bite 34 (Researcher Intervention):** On the participant’s request, we re-started the code with the “vertical skewer” motion primitive hardcoded, as that acquisition action tends to have better success with strawberries.
- **Bite 44 (User-resolved Off-nominal, Outlier Bite Duration):** The motion from the user’s mouth failed soon after it started. The user hit “retry” and it completed smoothly.

L. CR2 Study Details

Fig. 18 shows the timeline of CR2’s meal. Although all other participants had their meals around a traditional lunchtime, CR2 had his meal between a traditional lunch and dinner time, resulting in him eating fewer bites before getting full. Notable events in this meal include:

- **Bite 0-1 (User-resolved Off-nominal):** The robot’s camera feed did not render on the app in bite selection. The user clicked the “reload video” button and then it did.
- **Bite 0-1 (User-resolved Off-nominal):** None of the masks returned from food segmentation aligned with the user’s desired bite. To address this, he re-invoked food segmentation with a new seed point.
- **Bite 1 (Bite Acquisition Failure):** The robot was off-center and therefore missed the bite (addressed in System Patch 2d).
- **Bite 1 (Researcher Intervention):** Researchers nudged the fork in the gripper, so it better aligns with the robot’s URDF model (addressed in System Patch 2d).
- **Bite 8 (Bite Acquisition Failures):** The fork didn’t go all the way down to the food, perhaps due to inaccurate depth readings near the edge of the camera view. Addressed in System Patch 2b.
- **Bite 9 (Bite Acquisition Failures):** The fork pushed the bagel piece to the side as it descended into the bagel,

because the curve of the bagel aligned with the curve of the fork tines.

- **Bite 14 (Bite Acquisition Failure):** The fork was off-center on the piece of broccoli (addressed in System Patch 2d).
- **Bite 15 (User-resolved Off-nominal):** Robot was unable to find a plan to move into the food. User clicked “back” and re-selected their desired food item. Partially addressed in System Patch 2g.
- **Bite 15-16 (Researcher Intervention):** Robot did an extremely large, multi-part motion when moving into the food. The participant clicked the emergency stop button, which immediately stopped the robot. Swivel issue addressed in System Patch 2c. Researchers restarted the software afterwards and the participant continued.

M. P4 Study Details

Fig. 19 shows the timeline of P4’s meal. Notable events include:

- **Bite 1-2, 2, 7 (User-resolved Off-nominal):** The force-torque sensor disconnected from WiFi (causing all robot motion to immediately stop), but reconnected shortly thereafter without researcher intervention.
- **Bite 1-2 (User-resolved Off-nominal):** None of the masks returned from food segmentation aligned with the user’s desired bite. To address this, they re-invoked food segmentation with a new seed point.
- **Bite 1-2 (User-resolved Off-nominal):** When the user switched between apps on his phone, the robot web app rendered smaller than expected. He resolved this by reloading the page.
- **Bite 1-2 (Researcher Intervention):** Since the plate location in the camera feed was too small for the user to click, researchers moved the plate as the participant zoomed into the image, ensuring the full plate was visible zoomed in.
- **Bite 1-2, 2, 3, 3-4, 7, 16, 21, 25-26 (User-resolved Off-nominal):** A full-screen “sign in to Google” pop-up opened on the participant’s browser. In two of those occasions, that caused the robot action to immediately be canceled by the web app (since it was no longer foregrounded). In all other occasions, the robot arm was already stationary, but nevertheless this off-nominal prevented the user from interacting with the system (akin to if they receive a phone call while eating). In all cases, the user closed the pop-up, clicked “resume” or “back” on the web app if robot action had been terminated, and continued his meal.




Fig. 19. The full timeline of P4’s meal.




Fig. 20. The full timeline of P5’s meal.

- **Bite 4, 13, 14, 26 (Bite Acquisition Failure):** The fork pushed the cauliflower to the side as it descended into it, because the curve of the cauliflower top aligned with the curve of the fork tines.
- **Bite 5 (Researcher Intervention):** Due to the force-torque sensor frequently losing WiFi connection, researchers raised the backpack containing the router up off of the ground.
- **Bite 13, 24, 28, 34 (User-resolved Off-nominal):** The browser’s “History” tab opened full-screen, perhaps triggered by the participant’s voice control due to the conversation. Details and resolution are the same as the “sign in to Google” pop-up above.
- **Bite 16 (User-resolved Off-nominal):** The robot’s code was hanging temporarily. The user used the app to pause, press back, and retry the action, which then succeeded.
- **Bite 20 (User-resolved Off-nominal):** The browser’s “Option” tab opened full-screen, perhaps triggered by the participant’s voice control due to the conversation. Details and resolution are the same as the “sign in to Google” pop-up above.
- **Bite 21 (User-resolved Off-nominal, Outlier Bite Duration):** The robot encountered an error moving from the user’s mouth (perhaps due to a phantom obstacle in the Octomap). The user left the robot there for a while as we were conversing and a researcher was serving him more food. Eventually the user hit “retry” and it resumed as expected.
- **Bite 22 (Researcher Intervention):** The bite acquisition ended in such a position where the only plan that could be found in the allotted time limit to move to the “resting” configuration involved a big swivel, which the threshold implemented in System Patch 2c was rejecting. Thus, a researcher terminated the code, increased the threshold, and restarted it.
- **Bite 29-30, 30 (User-resolved Off-nominal):** The “live video” view of the app automatically popped up (perhaps

the voice control running on the user’s phone mistakenly heard a command and opened it). Both times, the user closed it so they could return to the main app screen and continue.

N. P5 Study Details

Fig. 20 shows the timeline of P5’s meal. Notable events include:

- **Bite 3 (Outlier Bite Duration):** The user forgot they have to tap a button to get the robot to move back from their mouth, and therefore left it at their mouth for around a minute as we were conversing, before remembering to click the button.
- **Bite 8, 10, 17 (Bite Acquisition Failures):** The robot arm went between pieces of mac, acquiring nothing. This is also partly because those pieces of mac were oriented with the “hole side up,” making it harder to skewer.
- **Bite 11 (Researcher Intervention):** Participant mistakenly pushed the emergency stop button. To address this, researchers manually restarted the code.
- **Bite 13 (User-resolved Off-nominal):** None of the masks returned from food segmentation aligned with the user’s desired bite. To address this, the user re-invoked food segmentation with a new seed point.
- **Bite 15, 24 (Bite Acquisition Failure):** The robot tilted the piece of food as it descended into it; thus, the fork tines did not skewer the food.
- **Bite 17 (User-resolved Off-nominal):** User mistakenly initiated bite transfer when the robot hadn’t acquired anything. On the “detecting face” screen, they clicked “move above plate” to have the robot return.
- **Bite 18, 28 (User-resolved Off-nominal):** Robot action encountered an error. The user clicked “retry,” and it proceeded smoothly.
- **Bite 22 (Bite Acquisition Failure):** The robot was off-center and only acquired a tiny piece of fish. After this

failure, the participant decided to acquire another piece of fish while that small piece was still on, which succeeded.

- **Bite 28 (User-resolved Off-nominal):** The robot action was stalling and/or hanging. The user clicked “pause” and “back,” which addressed it.

APPENDIX E

STUDY 2: SINGLE-USER, IN-HOME DEPLOYMENT

This section contains additional details of Study 2, beyond the core details presented in Sec. V.

A. Caregiver Demographics

Table XI shows the demographics of the three caregivers who were present during the deployment. Although all worked with multiple care recipients, their familiarity with assistive technology came from the assistive technologies that CR2 used. Thus, they were all familiar with Alexa, voice control, mouth joystick, power wheelchairs, a hospital bed, ceiling lift, accessible van, and more. The questions “How familiar are you with assistive technology for people with motor impairments” and “How familiar are you with robots” were each on 5-point Likert scales: “Not at all familiar,” “Slightly familiar,” “Somewhat familiar,” “Moderately familiar,” and “Extremely familiar.”

B. Study 2 Schedule & Overview

Fig. 21 shows the meal schedule for Study 2. Of the 5 consecutive days, 3 were wheelchair days and 2 were bed days. On wheelchair days, CR2 used the robot to feed himself breakfast and dinner. On bed days, he used the robot to feed himself snack and dinner³⁴

We co-decided the meals with CR2, informed by his preferences and the robot’s capabilities. We decided on the first 5 meals before the deployment began and on the latter 5 meals during the deployment week. The Wednesday breakfast (avocado toast) and part of the Friday dinner (roasted carrots and zucchini) were made by C2 and C3, respectively; all other meals were purchased from local stores and restaurants.

All meals but the Tuesday snack had a caregiver who was scheduled to be there for the entire meal. As a live-in caregiver, C2 popped into several of the meals for part of the time (typically the latter half): this occurred on the Mon dinner, Tues snack and dinner, and Thurs snack and dinner.

Over the course of the deployment, CR2 had the robot on his right, left, and front side. Across wheelchair days, he tried all three of the robot’s mounts: wheelchair, hospital table, and tripod.

Over the course of the deployment, CR2 accessed the web app using his laptop and his phone. Regardless of the device, he used the mouth joystick to control the device.

³⁴Bed-days have more required activities of care in the morning. Since they are more rushed, CR2 opted to not use the robot to feed himself breakfast on those days.

C. Semi-Structured Interview Questions

The exact questions we asked CR2 during the semi-structured interviews varied based on the flow of the conversation. Below are a superset of questions we asked to start conversations (conversation-specific follow-up questions not included):

1) Questions for the Care Recipient:

a) After Every Meal:

- What was the experience like of using the robot to eat this meal [while doing the co-occurring activity]?
- What challenges did the [co-occurring activity] introduce? How did you overcome those challenges using the robot? How could the robot be improved to better address these challenges?
- What surprised you about using the robot to eat this meal [while doing the co-occurring activity]? What would you do differently if you were to use the robot to eat this meal [while doing the co-occurring activity] in the future?
- As it stands right now, can you envision using the robot regularly to eat this meal [while doing the co-occurring activity]? If not, what would need to be changed for you to envision yourself using it regularly?
- What changes in your environment or norms, if any, would you be willing to do for this robot to work?

b) After The Deployment:

- Reflecting on this week, what went well? What went poorly? What surprised you?
- Think about your meal routine this week compared to your meal routine in past weeks. What aspects of the meal routine this week did you prefer compared to your meal routine in past weeks? What aspects of the meal routine in past weeks did you prefer compared to your meal routine this week?
- For each of the following contexts, what went well and what went poorly about eating a meal with the robot?
 - **Location:** In-bed vs. wheelchair
 - **Time:** Breakfast, Snack, Dinner
 - **Co-occurring activity:** working, watching a movie, conversing, while a caregiver does another activity of care
 - In which of these contexts would you like to continue using the robot-assisted feeding system? In which would you prefer being fed by a caregiver?
 - Do more contexts come to mind in which you’d like to try using the robot-assisted feeding system?
- Let’s walk step-by-step through each part of the robot-assisted feeding system. Please share any reflections or feedback you have on those components, both for the web app and the robot.
 - Customization
 - Bite selection
 - Bite acquisition
 - Bite transfer
 - Auto-Continue
- Is there anything else you would like to share with us?

ID	Age Group	Gender	Years as Caregiver	Years Worked with CR2	Live-in?	Impairments of Care Recipients	Num Deployment Meals	Familiarity with Assistive Technology	Familiarity with Robots
C1	25–34	F	0.5	0.5	✗	SCI, muscular dystrophy	3	“Somewhat familiar”	“Slightly familiar”
C2	55–64	M	25	25	✓	people with motor impairments	8	“Extremely familiar”	“Somewhat familiar”
C3	35–44	F	7	7	✗	people with motor impairments	4	“Extremely familiar”	“Not at all familiar”

TABLE XI
CAREGIVER DEMOGRAPHICS FOR STUDY 2.

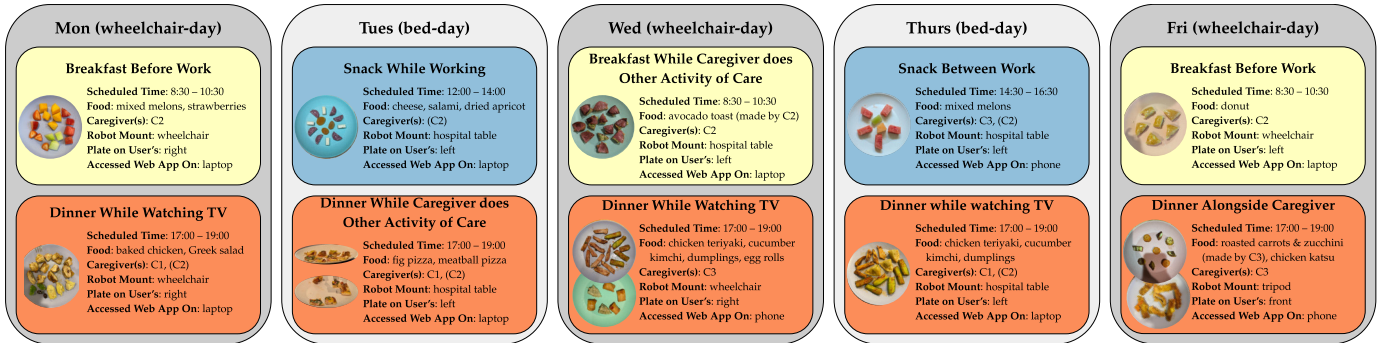


Fig. 21. An overview of the deployment schedule. Caregivers with names in parentheses were there for part, not all, of the meal.

2) Questions for Caregivers:

- Based on what you’ve seen of the robot arm, what do you think are the benefits of a robot-assisted feeding system? What are the drawbacks?
- What might change in your caregiving routine if CR2 had access to a robot-assisted feeding system?
- Would you feel comfortable working in a house where the person feeds himself with a robot-assisted feeding system? Why or why not?
- What do you think of the setup procedure for the robot (explain it if need be). Would you be willing to set up the robot for CR2? What type of setup procedure would you like? How can we make the setup procedure simpler for you?

- How do you think such a system would impact CR2’s health and well-being?
- Is there anything else that you would like to share with us?

D. Data Analysis

Initial transcription of all quotes was done by OpenAI’s Whisper speech recognition model³⁵. Subsequently, one researcher listened to all the video recordings and corrected mis-transcribed participant quotes, both during the meal and after the meal. That researcher then used thematic analysis [103] to tag quotes with their key themes.

³⁵<https://apps.apple.com/us/app/whisper-transcription/id1668083311?mt=12>