Achieving Deployable Autonomy in Robot-assisted Feeding for People with Motor Impairments

Amal Nanavati

Dec 8, 2023

Paul G. Allen School of Computer Science & Engineering
University of Washington
Seattle, WA 98105

Thesis Committee:
Maya Cakmak, University of Washington (Co-Chair)
Siddhartha S. Srinivasa, University of Washington (Co-Chair)
Anat Caspi, University of Washington
Henny Admoni, Carnegie Mellon University
Kim Ingraham, University of Washington (GSR)

Abstract

Over 1 billion people worldwide are estimated to experience significant disability, which impacts their ability to independently conduct activities of daily living (ADLs) such as eating, ambulating, and dressing. Personal physically assistive robots have emerged as a promising technology to help users conduct ADLs, thereby restoring independence and reducing caregiver burden. However, despite decades of research on personal physically assistive robots, deployments of them outside the lab are still few. A big reason for this is that every user is unique: their requirements for the robot, their needs, and their contexts of use are all unique. Thus, in order to be adopted, the robot has to work well for each unique user in their context(s) of use.

In this proposal, I focus on robot-assisted feeding as a case study for how we can achieve deployable autonomy in personal physically assistive robots. Our ultimate goal is to enable a robot-assisted feeding system to feed any user, in any environment, a meal of their choice, without researcher intervention, and in a way that aligns with their preferences. Towards that goal, I present completed and proposed work focused on the following themes: (1) understanding users' needs and priorities when it comes to robot-assisted feeding; (2) generalizing the robot's bite acquisition to the variety of foods users may want it to feed them; (3) empowering the user to customize the robot-assisted feeding system to their needs and environment; and (4) designing a system that can be used without researcher intervention. The former two works are completed, and the latter two are proposed and ongoing. Throughout this research agenda, I intend to conduct evaluations of the robot-assisted feeding system in environments outside the lab—e.g., conference rooms, cafeterias, home environments—with a diverse set of users, thereby assessing how close we are to the aforementioned ultimate goal.

Contents

1
Introduction

According to estimates by the World Health Organization (WHO), 1.3 billion people worldwide experience significant disability [132]. These disabilities threaten one's ability to independently perform activities of daily living (ADLs) such as eating, ambulating, and dressing [34], leaving them reliant on a caregiver for assistance completing such activities. Because most people with disabilities wish to live independently in their home [37; 49], research for several decades has focused on developing personal physically assistive robots to help people with disabilities independently perform ADLs [127]. But despite considerable advances, deployments of such robots are few [90], with notable exceptions including [46; 47; 64; 119; 93; 61].

Figure 1.1: Our robot-assisted feeding system being used in a crowded, dynamic, and loud social setting.

What makes deploying state-of-the-art personal physically assistive robots so difficult? One reason is that every user's needs, preferences, and environments are unique, which requires the robot to be customizable and generalizable. Another reason is that the real-world contexts these robots would be deployed in differ significantly from the lab contexts they are developed in (Fig. 1.1), giving rise to a variety of off-nominal scenarios the robots has not been designed for.

Figure 1.2: A community researcher eating a piece of chicken tenders from the robot-assisted feeding system.

Consider the case of a robot-assisted feeding system for a user with motor impairments (Fig. 1.2). First, the user would likely want to eat a variety of food items—sandwiches, noodles, salads, etc. Thus, in order to be deployed, a robot-assisted feeding system must be able to **generalize across a large variety of food types**. Second, the user will likely have impairment-specific needs; for example, not being able to chew with one side of their mouth, or not being able to move their neck beyond a certain distance. Thus, in order to be deployed, a robot-assisted feeding system must be able to **customize to its user's needs and preferences**. Third, the user will likely use it in a variety of environments; for example, some users have oversized wheelchairs, so while they are able to fit under a table at home, at restaurants they have to turn perpendicular to the table, leaving the plate on their side. Thus, in order to be deployed, a robot-assisted feeding system must be able to **customize to the environments its user eats in**. Finally, a variety of off-nominal scenarios might arise over the course of a meal: the caregiver might move the plate when serving the user more food; the user might need to take a break from the meal to cough; the robot might not detect the user's face based on how they are sitting; etc. Thus, in order to be deployed, a robot-assisted feeding system must be **robust to the variety of off-nominal scenarios that will inevitably arise in any deployment**.

1.1 Proposed Research Agenda

In this work, I propose a research agenda where we use robot-assisted feeding as a case study to investigate how we can achieve deployable autonomy in personal physically assistive robots. Robot-assisted feeding serves as a good case study for two reasons. First, a robot arm designed for people with motor impairments to teleoperate already exists and has adoption[1], which can

[1] Kinova® Jaco® Arm

form the hardware platform that we develop a robot-assisted feeding system on top of. Second, multiple commercial robot-assisted feeding systems have been developed over the years [4; 3; 120; 9; 104; 1; 76; 54]. While they have not yet reached widespread adoption, they demonstrate that there is a market for robot-assisted feeding systems[2]. Thus, we believe robot-assisted feeding systems have a path to translation from research to an actual product, which makes it a good case study for achieving deployable autonomy in personal physically assistive robots.

[2] This market is of considerable size; as of 2010, 1.8 million people in America alone needed assistance eating [125].

The proposed research agenda is guided by the following research question:

RQ-Thesis *How can we develop a deployable robot-assisted feeding system that can feed any user, in any environment, a meal of their choice, while aligning with their preferences?*

In order to answer RQ-Thesis, I propose investigating the follow intermediate research questions:

RQ1 What challenges do people with motor impairments face during dining, and how should a robot-assisted feeding system be designed to address those challenges?

RQ2 How can a robot-assisted feeding system feed users the large variety of food items they want to eat?

RQ3 How can a robot-assisted feeding system customize to users' needs and environments?

RQ4 How can we take a functional robot-assisted feeding system and make it deployable?

Figure 1.3: List of research questions that seek to ultimately answer RQ-Thesis.

Figure 1.4: The robot-assisted feeding system consists of a wheelchair-mounted robot arm with an on-board RGB-D and force-torque sensor, router, and compute.

RQ1 is the formative research that underlies the entire research agenda, enabling us to understand and be guided by participant needs and preferences. RQ2 focuses on feeding users "a meal of their choice." RQ3 focuses on feeding "any user, in any environment...in a way that aligns with their preferences." RQ4 focuses on "develop[ing] a deployable robot-assisted feeding system." RQ1 and RQ2 are completed [89; 42], while RQ3 and RQ4 are proposed ongoing work.

To demonstrating our progress towards RQ-Thesis, I propose multiple deployments of the robot-assisted feeding system. Most will be single-meal deployments, where the system feeds a participant with motor impairments an entire meal in an out-of-lab eating environment—e.g., a cafeteria, atrium, or conference room. One will be a *week-long deployment of the system in a participant's home*. From each deployment, we will gather quantitative and qualitative data on the robot's performance and participants' experiences, and use it to assess steps for future work.

1.2 *Approach: Community-Based Participatory Research*

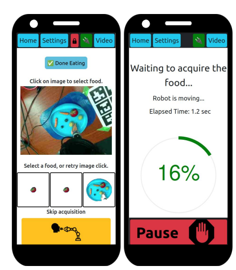

Figure 1.5: The system also consists of a web app for the user to start and stop robot motion, provide inputs such as the bite they want to eat, and customize their feeding experience.

The entire proposed research agenda follows principles from community-based participatory research (CBPR). CBPR is a method where academics researchers work equitably with community members throughout all research stages, from ideation to dissemination [48; 130]. CBPR is rooted in the belief that community members and academic researchers each bring unique skills, expertise, and lived experiences to the team; for example, while academic researchers are familiar with rigorous research methodologies, community members are familiar with the nuances of the problem, through lived experiences. Thus, addressing a community need requires an equitable partnership between academics and the community, which involves sharing power, resources, credit, results, and knowledge [102; 85]. CBPR has been used in the health sciences for decades [59; 130], and is increasingly used in assistive technology research [89; 81; 27; 71; 15].

All completed and ongoing work in this research agenda has been conducted in close collaboration with two community researchers. These individuals helped define the research questions, run the studies, design the robot-assisted feeding system, disseminate the results, and more.

1.3 *Robot-Assisted Feeding System Overview*

Hardware: The robot-assisted feeding system we use in this work has been built over the years by multiple generations of Personal Robotics Lab members, including myself. This fully portable system (Fig. 1.4) consists of a 6-DoF robot arm mounted to a power wheelchair. The arm holds a fork and has an on-board RGB-D camera and force-torque sensor. The camera is used to detect bites of food and the user's face. The force-torque sensor is used to determine when the fork has contacted food and as a crucial safety measure to ensure the robot stops moving as soon as it senses an unexpected force. The system also consists of a laptop and router, both mounted on the back of the wheelchair. All system components draw power from the wheelchair, resulting in no external cables required to run the system.

Software: The user interacts with the system through a web app (Fig. 1.5), that was co-designed with the community researcher following the principles in "Ability-Based Design" [131]. The web app invokes "actions," or modular units of computation, on the robot arm. One such action is "AcquireFood," which takes as input a detected food item and moves the robot to acquire it with its fork. Another such action is "MoveToStagingConfiguration," which moves the robot arm to a fixed "staging" configuration from which the user's face should be visible. Other actions are "DetectFace" and "MoveToMouth," which are self-explanatory. In fact, the whole robot-assisted feeding can be considered a state machine where one or more "actions" move the robot arm between key stationary arm configurations or waypoints, shown in Fig. 1.6.

Under-the-surface, each "action" is a behavior tree [33] that computes the goal of the motion, plans a collision-free path using RRT* [63] or a cartesian interpolator, and executes that path with a joint trajectory controller [31]. Collisions are represented through simulated static objects—a table, wheelchair, and expansive bubble around where the user would sit—and a dynamic map with real-world obstacles perceived by the depth camera [56].

In addition to using the web app to invoke robot actions, the user can also use it to provide inputs to the system (e.g., which bite they want), stop and restart robot motion, and customize the robot. Users are able to use the web app with any assistive technology that they use to interact with smartphones, tablets, and/or computers.

Additional system details, particularly focused on deployability, are presented in Ch. 6 (RQ4).

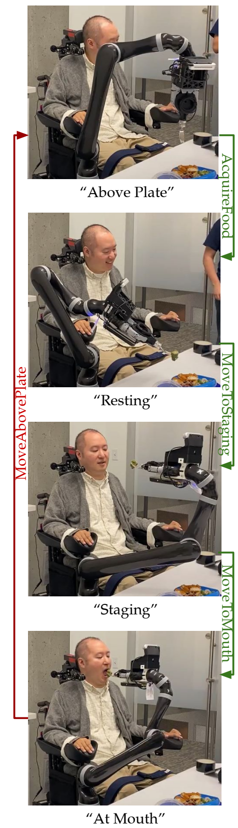

Figure 1.6: Key waypoints where the robot waits. "Above Plate," "Resting," and "Staging," are fixed, whereas "At Mouth" depends on the perceived mouth location. Green arrows (right) shows progression when feeding the user, and red (left) when moving back for another bite.

1.4 *Roadmap*

The remained of this document is structured as follows. Ch. 2 discusses key related works in robot-assisted feeding and personal physically assistive robots. Ch. 3 presents the results of our investigation into people with motor impairments' needs and preferences for a robot-assisted feeding system (RQ1). Ch. 4 presents the results of our investigation into generalizing the robot's bite acquisition (RQ2). Ch. 5 presents proposed work to make the system customizable to users' needs and environments (RQ3). Ch. 6 presents ongoing work to make the system deployable and my proposed evaluations of the system (RQ4). Finally, Ch. 7 presents the proposed evaluations and timeline of this research agenda.

2
Related Works

2.1 A History of Robot-assisted Feeding

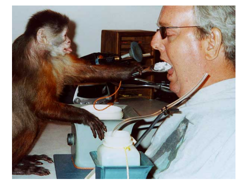

Figure 2.1: A trained monkey feeds a person with motor impairments. Reprinted from Envisioning Access.

Enabling people with motor impairments to eat independently has been a goal of research for around 50 years, and has been covered in the following surveys [104; 124; 70; 2; 19; 91].

One of the early research directions in the 1970s involved training capuchin monkeys as service animals to feed people with motor impairments [78] (Fig. 2.1), an effort that continued for decades [53]. Also in the 1970s, an early robot-assisted feeding system, the Morewood Spoon Lifter (Fig. 2.2), was developed. This portable system involved strapping a metal rod around the user's head, which they used to shovel food into the spoon and to press a switch that moved the spoon from table-level to mouth-level. The US Department of Veterans Affairs (VA) clinically evaluated this device with 16 people with quadriplegia and studied it during a 3 year home deployment with one user [104]. This device was later renamed the "Winsford Feeder" (Fig. 2.3), which was manufactured and sold as a commercial product until at least the 2010s [104; 91].

Figure 2.2: The Morewood Spoon Lifter, developed in 1974. Reprinted from Appendix C of [104].

While the "Winsford Feeder" focusing on being portable, other systems focused on being multi-purpose. Developed in the 1980s, the "Robot Arm Work Table" consisted of a desk with a fixed robot arm and other tools mounted to it. The system was designed to enable people who only have neck mobility to drive their wheelchair up to it and use the robot arm for a variety of different tasks, including picking and hanging up the phone, typing on a computer, and eating from a bowl. This system was deployed and clinically evaluated by the VA with 20 people with quadriplegia, over up to a year, in environments as diverse as family home, a nursing home, and a hospital [114]. Around this time, other stationary, multi-purpose robot-assisted feeding systems were also being developed, such as the Handy 1 (Fig. 2.4), which was designed to help an 11 year old boy with cerebral palsy eat, drink, brush his teeth, and more [127].

Figure 2.3: The Winsford Feeder was sold until at least the 2010s. Reprinted from North Coast Medical.

By the 2000s, multiple commercial products for robot-assisted feeding were on the market. These include: Bestic, first developed in 2004 and sold in 2012 [76] (Fig. 2.5); Obi, first developed in 2009 and sold in 2016 [9] (Fig. 2.6); Neater Eater, in development in the 1990s [84] and sold in the early 2000s [54] (Fig. 2.7); My Spoon, in development in the 1990s and sold in 2002 [120]; the aforementioned Winsford Feeder; and more [2; 91]. These products are all table-mounted and have the robot execute fixed trajectories to acquire food and move it to the user's mouth. They have undergone considerable user testing, including clinical evaluations, which have shown positive results in terms of being able to eat a full plate of food [76], feeling more independent and confident [69; 76], and having improved posture [69]. Despite the positive results, these devices have struggled to achieve long-term adoption, with all but the Obi and Neater Eater being discontinued. Some of the shortcomings include: being unable to acquire users' desired food items or acquiring too little food [119; 66]; dropping food [77]; and requiring precise positioning of the user, sometimes resulting in strained muscles [94]. These shortcomings can be traced to

Figure 2.4: The Handy 1 was developed in 1987 to help a boy with cerebral palsy to independently eat. Reprinted from [127].

an inability to autonomously sense and react to the environment, e.g., adapting the acquisition strategy to perceived properties of the food, adapting the transfer motion to the face pose.

Contemporary research in robot-assisted feeding has largely sought to address the aforementioned shortcomings. Specifically, the systems currently being used in research typically have more sensors than commercial systems—including RGB cameras, depth cameras, and/or force-torque sensors—which they use to be more adaptive to the food, user, and environment. Modern research systems take a variety of forms, including wheelchair-mounted robot arms [40; 18; 97], table-mounted portable robot arms [118], table-mounted fixed robot arms [14], and mobile manipulator robots [100; 93]. Sec. 2.2 and 2.3 focus on these systems.

Figure 2.5: Bestic feeds its founder, Sten Hemmingsson. Reprinted from Bestic AB.

Figure 2.6: The Obi is a commercial device that has been sold in the US since 2016. Reprinted from MeetObi.

Figure 2.7: The Neater Eater is a commercial device that has been sold in the UK for around two decades. Reprinted from Neater Eater.

Figure 2.8: A caregiver feeds a care recipient in an assisted-living center. Reprinted from [16].

2.2 *Formative Studies in Robot-assisted Feeding*

A few formative works have studied participants needs and priorities when it comes to robot-assisted feeding. Bhattacharjee et al. [16] conducted a contextual inquiry in an assisted-living center, where they observed meals, showed care recipients and caregivers a video of robot-assisted feeding, and interviewed them (Fig. 2.8). They synthesized several evaluation indicators for robot-assisted feeding systems, including their technical function, technology robustness, information gaps, usability, social acceptance, and system integration. Pascher et al. [101] observed participants consuming a meal in their homes, had them experience robot-assisted feeding through a virtual reality headset, and interviewed them. They present several recommendations for the development of robot-assisted feeding systems, including that such systems should not involve lengthy familiarization periods, should be compact and unobtrusive, and should allow users to perform daily tasks being feeding. Finally, Kim et al. [65] conducted focus group interviews with care recipients, caregivers, and doctors to understand their priorities for robot-assisted feeding. One of their key findings was that each group has different preferences; for example, doctors and caregivers want the robot to move twice as slow as care recipients do.

The work presented in RQ1 [89] adds to this space of formative studies by focusing on user's needs and priorities when it comes to robot-assisted feeding in *social settings*.

2.3 *Technical Advances in Robot-assisted Feeding*

Technical research often divides robot-assisted feeding into two main sub-tasks: bite acquisition (Fig. 2.9) and bite transfer (Fig. 2.10). Within acquisition, prior works have focused on the robot's ability to acquire food with a fork [43; 52; 122; 121], spoon [110; 95; 58], chopsticks [133; 96], or multiple tools [119; 100; 45]. Some of these works focus on the motion primitives a robot should use to acquire food [52; 17; 40; 121], others focus on learning which motion primitives to use on novel food items [43; 44], and yet others focus on chaining together motion primitives for more complex acquisition (e.g., pushing food together before scooping it) [122]. The work presented in RQ2 [42] adds to this by presenting a pipeline to learn motion primitives from human data.

Other works have focused on bite transfer. These works have investigated planning motions that account for user comfort [13], learning from caregiver demonstrations [23], and revealing the coupling between how a bite is acquired and how it can be transferred to the mouth [40]. Although most works have the robot arm stop a few centimeters in front of the user's mouth, some recent works have investigated in-mouth bite transfer [116; 61] (Fig. 2.11).

Finally, some works have extended beyond bite acquisition and transfer. Some focus on predicting what food users want [60], others on detecting food [52; 38] and the user's mouth [103], others on predicting when users want a bite [97; 52], and others on detecting anomalies during feeding [99]. Recent works have developed a simulation environments for caregiving tasks [134; 80].

2.4 *Deployments of Physically Assistive Robots*

Beyond robot-assisted feeding, we may be able to draw relevant insights from the broader field of personal physically assistive robots for people with disabilities. The state-of-the-art in physically assistive robotics research has been presented in multiple survey papers, including [90; 86; 98; 32]. In this section, we focus on works that deployed their physically assistive robots, since in-the-wild deployments provide important research insights [62] and have proved crucial to translating past robot-assisted feeding systems from research to product (Sec. 2.1).

Figure 2.9: Bite acquisition is the process of acquiring a bite of food. Reprinted from [38].

Figure 2.10: Bite transfer is the process of moving a bite of food to the user's mouth. Reprinted from [40].

Within the realm of robot-assisted navigation, one work allowed blind visitors to a museum to use a robot and associated app to freely explore the museum for 1.5 hours [64]. Another gave a robotic walker to a blind user, allowing them to freely use it, without researcher intervention, for 2 months [47]. Within the realm of assistive teleoperation, one work demonstrated that a person with quadriplegia was able to teleoperate a robot in his house to shave himself [51]. Another work deployed a robot in a user's home for 4 weeks, with researcher presence, and found that he was able to use a teleoperation interface to feed himself, scratch an itch, and play cards [93]. Within the realm of robot-assisted feeding, one study deployed a table-mounted feeding robot in a user's home for 5 days [119], whereas another had a wheelchair-mounted robotic system feed a participant one meal in their home [61] (Fig. 2.11). Notably, there is space for longer deployments and deployments where users freely use their wheelchair-mounted robot-assisted feeding system without researcher intervention. This motivated RQ-Thesis.

Figure 2.11: A robot-assisted feeding system feeding a user in their home. Reprinted from [61].

A common technique to achieving robust robot deployments in human-robot interaction (HRI) is relying on humans-in-the-loop. This includes a mobile robot asking bystanders for help to press the elevator button [112], asking dedicated helpers to put it back on its charger [126], or asking for suggestions on how to improve at its task [74]. This insight informed our approach to **??**.

2.5 *Customization in Physically Assistive Robots*

Several of the aforementioned formative studies (Sec. 2.2) found that users have a variety of needs, preferences, and contexts of use, all of which the robot must be able to customize to. In fact, Nickelsen [94] found that users and caregivers inevitably tinker with assistive technology to get it to work for them and, if unable to, stop using the technology. However, despite the importance end users and caregivers place on customizability, considerably fewer research papers focus on providing customizability, either in robot-assisted feeding system or in personal physically assistive robots in general.

Of the systems that do allow customization, some focus on customizing the user interface [107; 105; 10; 30]. Others focused on allowing users to customize the level of autonomy the robot exhibited [135; 108]. Yet others focused on enabling users to customize task-specific functionality, such as Canal et al. [25] which found that users' satisfaction with robot-assisted feeding increased when they were able to customize how close the robot got to their mouth and its speed (Fig. 2.12). In terms of when customization takes place, some works allow the user to customize the system during execution [135; 105], while others have a calibration phase before execution where users customize the system [10; 30; 8; 25].

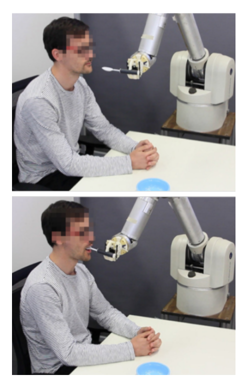

Figure 2.12: Canal et al. allowed users to customize the robot's speed and how close it got to their mouth. Reprinted from [25]

Most existing works focus on customizing single-dimensional parameters. To the best of our knowledge, RQ3 is the first that focuses on users customizing entire arm configurations, which is a 6-dimensional parameter that impacts the motion the robot-assisted feeding system takes.

3
Understanding Users' Needs (RQ1)

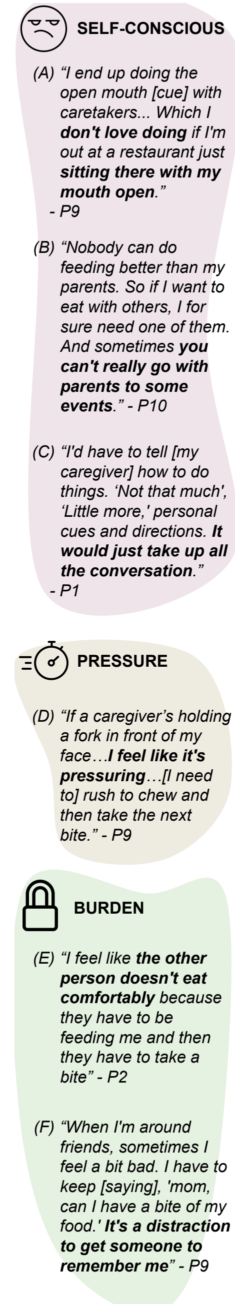

Figure 3.1: Negative emotions participants felt during social dining. Full figure in [89].

This chapter presents the first research foci: conducting formative research to understand the problem and design space for robot-assisted feeding. Specifically, we focus on the following:

RQ1 *What challenges do people with motor impairments face during dining, and how should a robot-assisted feeding system be designed to address those challenges?*

Although there is prior formative research in robot-assisted feeding (Sec. 2.2), one aspect that has been under-studied is social dining. Social dining has been shown to have biological, psychological, and cultural benefits for those who engage in it [73; 82; 128; 28; 41; 115]. Unfortunately, people who rely on caregivers to eat are often excluded from the benefits of social dining, with shared meals being less about *socialization* and more about *functionality* (e.g., meal prep, food intake) [87; 79]. While some summative studies have touched upon specific aspects of the social dining, such as when to feed the user [97], to the best of our knowledge there has not been a comprehensive formative study into users' needs and priorities for robot-assisted social dining. To address this gap, **we approach RQ1 from the perspective of social dining**, by focusing on how participants might use their robot-assisted feeding system in social settings, and on how their feeding experiences differ between individual and social contexts.

3.1 Methods

We interviewed $n = 10$ participants, primarily from the community researcher's connections. The inclusion criteria were to have a permanent motor impairment and rely on a caregiver to be fed.

The interview consisted of two stages. First, we discussed participants' current social dining routines. Second, we watched videos[1] (Fig. 1.1) that showcased various features of the robot-assisted feeding system in social settings and discussed participants' preferences over those features. The community researcher helped design the videos and lead the interviews.

We used thematic analysis [117] to analyze video recordings from the design interviews. To develop the codes and themes that emerged from the data, two researchers independently coded each interview recording and performed calibration exercises to ensure consistency [68]. The full publication [89] and supplementary materials [6] contain more details on the methods: participant demographics, interview questions, video descriptions, etc.

[1] https://youtube.com/playlist?list=PLvoSEVdRS7GqvB1eWGUrEvMwfNgdcbuMt

3.2 Results

Participants faced **repeated challenging experiences** that led them to **prefer not to eat socially**. *"I don't like it. I'll arrive and be like 'nope, I'm good, I already ate'... A lot of people eat out for enjoyment. For me, eating is a necessity, I don't do it for fun."* (P1). Fig 3.1 provides an overview of these challenges.

One challenge participants experienced was **caregivers' lack of consistency in meeting their needs**, because different caregivers feed them differently. Some feed too fast, causing **pressure** (Fig 3.1D), while others feed too slowly, causing **frustration**. Some offer bites that are too large, a **choking hazard**, while others' are too small: *"One day my dad's shoving half a chicken down my throat, the next a nurse is cutting the tiniest pieces; I'm like: 'I'm gonna be here for centuries!"'* (CR). To cope, some participants rely on a few consistent caregivers to feed them but felt that it can be inappropriate to bring specific caregivers to some social events (Fig 3.1B). Importantly, users seek both **customization** and **consistency** in their feeding experience; while select caregivers can provide a customized feeding experience, the reality of multiple caregivers precludes then from having a consistently customized experience. A customizable robot-assisted feeding system, which we focus on in RQ3, can achieve both.

Participants also faced challenges due to **a mismatch between environmental factors and their needs**. One user needed to tilt his wheelchair to regulate blood pressure and was constantly concerned: *"Am I going to... tilt back and crash into a waiter?"* Others adjust how they sit, making it **difficult to interact with others**: *"My chair is oversized, so I don't fit going straight into a table. I have to sit sideways."* (P4). This reveals the importance of the robot-assisted feeding system working in a variety of environments, which we also address through customizability in RQ3.

3.2.1 *Design Principles*

Across participants, several common themes emerged as to why they prefer certain robot features over others. These themes (Fig. 3.2) should guide the design of robot-assisted feeding systems.

Participants wanted a robot-assisted feeding system that works **reliably**, without errors. This includes errors feeding the user—such as dropping the food (Fig 3.2A) or colliding with the user (Fig 3.2D)—as well as errors given the social context—such as blocking the user's view of their conversational partner (Fig 3.2C) or infringing on their neighbor's personal space (Fig 3.2B). Note, however, that although participants expected the robot to complete its tasks reliably, they did not expect it to do every task involved in feeding them a meal. For example, participants were willing to have a caregiver cut their food into bite-sized pieces and then have the robot feed them: *"I haven't a problem having my wife cut up my food and I say 'give me a bite' and [the robot] picks up whatever is available. The problem that we are solving with this is relying on someone sitting next to me for the entire meal giving me spoonfuls of food"* (P8).

Participants also want **control** of their robot; they want to decide when and what the robot feeds them (Fig 3.2E-F) and to have supervisory control to stop the robot at any time (Fig 3.2D). However, there are limits to how much control the user wants; notably, users did not want low-level control over the robot's bite acquisition motion, and those who had tried to teleoperate a robot to acquire a bite found that it took them anywhere from $5 - 40$ minutes for a single bite.

These insights on reliability and control inform **??**'s focus on the system working reliably without researcher intervention, and our approach to achieving it via human-in-the-loop control.

Finally, participants emphasized the importance of customizing the robot-assisted feeding system to everyone's needs (Fig 3.2G) as well as the different environments people eat in (Fig 3.2H). This motivates our focus on customization in RQ3.

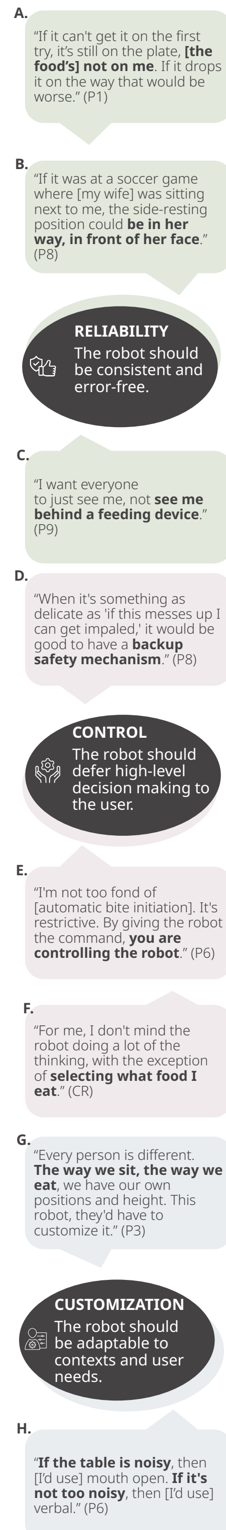

Figure 3.2: Three of the eight design principles for robot-assisted feeding. Full figure in [89].

4
Generalizing Bite Acquisition (RQ2)

This chapter presents the second research foci: expanding the space of food items the robot can acquire to cover the foods a user may want to eat. Specifically, we focus on the following:

RQ2 *How can a robot-assisted feeding system feed users the large variety of food items they want to eat?*

Given participants' willingeness to have a caregiver cut their food (Sec. 3.2.1), **we narrow our focus to bite-sized foods**. This also aligns with the robot-assisted feeding system (Sec. 1.3), which uses a single fork and can't cut many food items. We contribute a pipeline to learn bite acquisition actions (e.g., skewer, scoop, twirl, etc.) from human demonstrations. We then use that pipeline to learn to acquire the foods our community researcher eats in a week: bready items like bagels and pizza, heterogenous items like sandwiches rice and beans, gelatenous items like jello, and stringy items like noodles. Additional details can be found in the full publication [42].

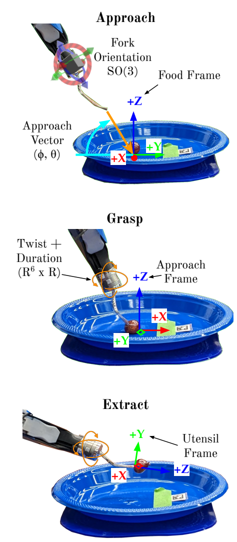

Figure 4.1: The action schema consists of an approach, grasp, and extraction motion. Robot motions are in orange. Reference frames are represented as three-color axes with X in red, Y in green, and Z in blue.

4.1 Acquisition Action Schema

Although previous work [17] qualitatively captured a taxonomy of human bite acquisition techniques with a fork (e.g., skewering, scooping, twirling, wiggling, etc.), they did not provide a quantitative representation that the robot can execute. Thus, we begin by defining a parameterized schema to represent acquisition actions. The design of this schema was guided by three goals: (a) the schema should cover the acquisition actions qualitatively described in [17]; (b) similar actions in the schema (e.g., skewering with vertical tines) should be close under some metric (e.g., Euclidean), to make it easier to discover novel actions; and (c) the motion should be agnostic to environmental factors like food orientation or table height.

Fig. 4.1 shows the schema, which is divided into three phases: **Approach**, **Grasp**, and **Extraction**. **Approach** encapsulates the initial motion into the food. Specifically, it consists of a straight-line motion from an initial fork pose to a target offset point within the food (e.g., in the center of the food item, 1cm below its surface). **Grasp** encapsulates the in-food manipulation, such as twirling noodles or scooping rice. Specifically, it consists of the robot executing a twist (linear and angular velocity) for a fixed duration or until it experiences a force or torque beyond a threshold. Finally, **Extraction** encapsulates the motion to lift the food out of the plate, and is also defined with a twist and duration. In total, this acquisition schema consists of 25 parameters (9 for acquisition, 9 for grasp, and 7 for extraction), the full details of which can be found in [42]. To ensure the action is agnostic to factors such as food orientation and table height, all motions are defined relative to the perceived bounding ellipsoid aroudn the food. Thus, even if the food is rotated or the plate is in a different position, the robot's executed action should still be the same relative to the food item.

4.2 *Gathering Human Data*

Although the 26 dimensional action schema is large, we hypothesized that only specific points (acquisition actions) within this schema will be commonly used to acquire food items. To identify those points, we had $n = 9$ able-bodied participants acquire a variety of food items and feed them to an actuated mouth. We selected 13 items based on the aforementioned foods the community researcher ate in a week. Each participant was given a plate with multiple bites of each food item, and requested to use a fork to acquire the item and move it to an actuated mouth. The fork had motion capture markers to track its position over time and the same force-torque sensor as the robot arm (Fig. 1.4) to gather haptic information. An RGB-D camera was also mounted above the plate. All-in-all, the dataset consists of 496 bite acqusions and can be accessed at [88].

4.3 *Learning Robot Acquisition Actions*

For each bite acquisition trial, we extracted an analogous motion within the acquisition schema. First, we computed the food frame by segmenting the bite that the fork first contacted. Second, we determined the timestamps delineating approach, grasp, and extraction by identifying peaks in the force and the robot's height above the table. Third, we computed the approach, grasp, and extraction parameters by linearly interpolating between waypoints between the respective timestamps. Finally, we ran k-medoids, sweeping k from 4-61, on the extracted actions. We selected $k = 11$ actions, as that was the elbow point of the curve.

4.4 *Evaluation*

We evaluated our system by selecting 14 food items—with 5 contained in and 9 not contained in the training dataset—and assessing the robot's success rate at acquiring the food items. The target success rate was 80%, since prior work found that usersare willing to tolerate up to 20% acquisition failures [18]. We conducted evaluations to investigate the following hypotheses:

1. **Coverage**: For each food item, at least one action in the set will achieve $\geq 80\%$ success.
2. **Learnability**: Using a state-of-the-art online learning framework [44], the robot will be able to converge to action(s) that have $\geq 80\%$ success within 11 attempts.

With regards to the first hypothesis, testing every action on every food item 10 times (1960 trials) revealed that for every food item but spinach there is at least one action that obtains $\geq 80\%$ success. Spinach, due to its thinness, fell off the fork more often, achieving 70% success. However, note that people often each spinach as part of a salad, where multiple leaves on top of each other might address the issue. The acquisition actions, which were learnt using unsupervised techniques, exhibited a variety of emergent behaviors that contributed to their success including in-food wiggling, tilted tines during extraction, and using vertical tines for high force. Finally, they performed equally to or better than the baseline actions from [38] on every food item; notably, the baseline actions were never able to acquire jello or spinach, whereas the learnt actions acquired them 70% and 100% of the time, respectively. This mostly proves the first hypothesis.

With regards to the second hypothesis, we ran rollouts of the LinUCB algorithm with post-hoc haptic context [44], where the robot started out with a uniform distribution over actions. Across food items, the robot converged to $\geq 80\%$ success rate within 8 rounds (e.g., 1 acquisition attempt each of 8 different pieces of that food). With each acquisition taking approximately 30 seconds, this reveals that the robot should be able to learn how to acquire a previously unseen food item with around 5 minutes of online training, for example through a pre-meal calibration phase. This proves the second hypothesis.

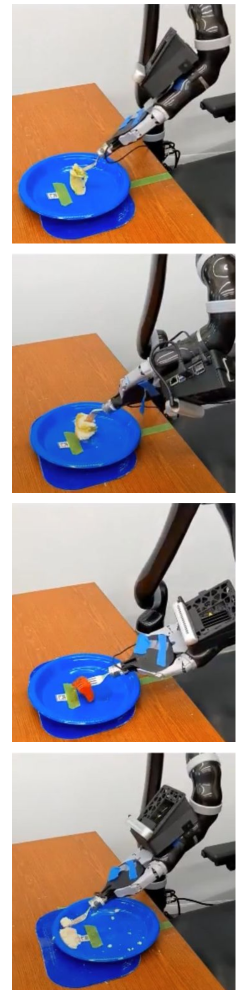

Figure 4.2: The learnt acquisition actions can acquire food items as diverse as lettuce, cut sandwich pieces, jello, and mashed potatoes (top-bottom) with at least 80% success rate.

5
Customizing to Users and Environments (RQ3)

This chapter presents the fourth research foci: empowering users to customize the system to their needs, preferences, and environments. Specifically, we focus on the following:

RQ3 How can a robot-assisted feeding system customize to users' needs and environments?

Without customizing to the user and their environment, the robot will not work for many users. If the robot does not get close enough to the user's mouth, from a direction they can chew from, it will be unusable. If the user has an oversized wheelchair and has to sit sideways relative to a table, but the robot expects the plate to always be in front of it, the system will be unusable. Canal et al. [24]'s taxonomy illustrates the wide variety of user preferences, and Table 5.1 presents customization goals specific to robot-assisted feeding (from RQ1 results).

A common approach to customization has the robot collect inputs—e.g., feedback, demonstrations—from users and learn the parameters of robot behavior that it thinks will align with those user preferences [20; 55; 106; 36]. We refer to this as **robot-driven customization**, because the mapping from user inputs to a robot parameter update is done by the robot. While these approaches have shown good performance in learning customized robot behavior, prior work has found that users feel frustrated when they are unable to change the robot behavior in the way they desire, or when there isn't sufficient transperancy into how their input will change robot behavior [7; 106; 21].

Our key insight is that users are the experts at what they want; by providing intuitive knobs to modify robot parameters, we can empower them to directly customize their robot-assisted feeding system. We refer to this as **user-driven customization**, because the user directly modifies the parameters of robot behavior. The key challenges of developing user-driven customization approaches are: (1) identifying robot behavior parameters that are expressive enough to capture diverse user preferences, while not being so extensive that they are unintuitive; and (2) providing sufficient transparency into parameters to enable users to make informed customization decisions.

In this chapter, I propose developing and comparing the tradeoffs between user-driven and robot-driven approaches to customizing the robot-assisted feeding system.

Table 5.1: A non-exhaustive list of user needs, preferences, and environmental factors that motivate customization.

Needs	Preferences	Environment
User can only move their head a certain distance to the fork	User wants the robot to (not) occupy their visual field	User's wheelchair is tilted relative to the table/plate
User must be fed from one side of their mouth	User wants the robot to take "natural" configurations	User is being fed in-bed, in a different relative position to the robot
User needs small bites to prevent choking	User wants the robot (not) to automatically move to their mouth	User wants the robot to not block the TV or social dining companion(s)
...	...	...

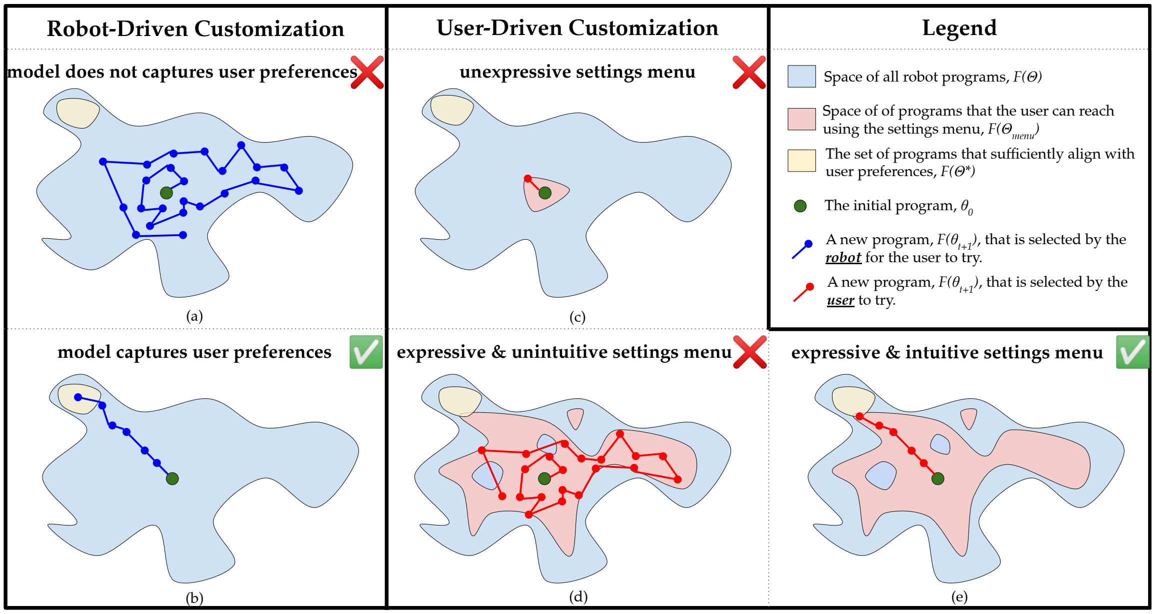

Figure 5.1: A visual representation of the framework for customizing robot behaviors (Sec. 5.1). (a) and (b) represents robot-driven customization where the model doesn't and does capture user preferences, respectively. (c) and (d) represents user-driven customization where the settings menu is unexpressive and expressive but unintuitive, respectively. (e) represents user-driven customization where the settings menu is both expressive and intuitive.

5.1 *Problem Formulation: Customizing Robot Behavior*

We formalize the problem of customizing robot behavior as a parameter search problem. Specifically, consider a parameterized family of programs that complete the feeding task $F(\theta) = f_\theta$. For each user and context of use, there will be some parameter setting(s) that sufficiently align with the user's preferences, e.g., $\Theta^* = \{\theta \in \Theta \mid h(f_\theta) \geq \eta\}$, where $h(f_\theta)$ represents how much the program f_θ aligns with the user's preferences, η is a minimum threshold of preference alignment that must be satisfied for the user to want to use the program. The problem of customizing a robot program involves finding a parameter $\theta^* \in \Theta^*$ that sufficiently aligns with user preferences.

Robot-driven customization typically requires an explicit model of user preferences, e.g., $h(f_\theta|w) = w \cdot \Phi(f_\theta)$ where the featureizer Φ extracts quantitative features from the program and the weights w represent the user's preference over those features. Robot-driven customization then seeks to learn the w that most aligns with user preferences, based on the the user's inputs (e.g., feedback or demonstrations). The benefit of this approach is that, by having an explicit model of user preferences, we can use standard optimization techniques to learn that w. The challenge is finding a featurizer Φ that represents the features users care about and a model h that represents how users care about those features. If a user's preferences are not captured in that explicit model, the system may be unable to find a program that aligns with their preferences.

User-driven customization gives the user access to a subset of parameters $\Theta_{\text{menu}} \subseteq \Theta$, for example through a settings menu, and allows them to freely tune parameters until they find one that aligns with their preferences, i.e., $\theta^* \in \Theta^* \cap \Theta_{\text{menu}}$. The benefit of user-driven customization is that we no longer need an explicit model of users' preferences; users can use any internal model to assess how well f_θ aligns with their preferences. The challenge is finding a set of parameters Θ_{menu} to expose to users that is: (a) expressive enough to capture diverse users' preferences, i.e., $\Theta^* \cap \Theta_{\text{menu}} \neq \emptyset$ across all users; and (b) intuitive enough to tune, i.e., when users experience a program f_{θ_t} they must be be able to identify a new parameter value θ_{t+1} such that $f_{\theta_{t+1}}$ better aligns with their (internal) model of preferences i.e., $h(f_{\theta_{t+1}}) > h(f_{\theta_t})$.

Fig. 5.1 illustrates this framework, as well as best- and worst-case scenarios of both robot-driven and user-driven customization[1]. The desirable scenarios in both robot-driven and user-driven customization, Fig. 5.1b and e respectively, are ones where the final program aligns with user preferences and the path to get there is short and convergent.

[1] This framework also allows us to represent shared customization, which is beyond the scope of this chapter.

5.2 What to Customize?

I propose focusing on two aspects of customization, that cover many scenarios in Table 5.1.

First, I propose allowing users to customize their **bite transfer** experience, which is crucial to ensure the user is able to eat the bite from the fork. Our robot-assisted feeding system's bite transfer is defined by two parameters: (a) the 6D "staging configuration" from which the user's face is visible (Fig. 1.6) and (b) the distance to the mouth the robot should stop at. A bite transfer is a straight line (cartesian) motion from the staging configuration to a point that far from the mouth. Thus, by customizing these two parameters, users can fully specify their desired bite transfer, based on whichever customization goal (e.g., in Table 5.1) is most relevant to them.

Second, I propose allowing users to customize the system's assumptions about the **relative positioning of the user, plate, and robot**. Two configurations are conditioned on these positions: the "staging" configuration from which the user's face must be visible, and the "above plate" configuration from which the plate must be visible (Sec. 1.3). Thus, allowing the user to customize these configurations should allow the system to work with a variety of relative positionings.

Given the importance of customizing the robot arm's key configurations, we will focus on that for the remainder of this chapter.

5.3 How to Customize?

I propose developing a user-driven and robot-driven method for customizing configurations, and evaluating their tradeoffs in a user study. Whereas the "above plate" configuration has only functional constraints (i.e., the plate must be visible), the "staging" configuration has both functional and user preference constraints (e.g., the face must be visible and the in-mouth motion must align with user preferences). Thus, I'll focus the user study on customizing the **staging configuration**.

5.3.1 User-Driven Customization

Nickelsen [94] found that users and their caregivers invariably tinker with their assistive robots in order to get it to work for their needs and environments, and stop using the technology if it that cannot be achieved. Thus, we focus on designing user-driven customization interfaces that facilitate such tinkering, by following the "Designing for Tinkerability" framework [109]. For example, the user-driven customization interface will provide "immediate feedback" by allowing users to try the robot motions that result from their parameter choice(s). It will also support "fluid experimentation" by giving users intuitive control to search through the parameter space and transparency into the technical constraints that inform parameter selection.

Specifically, to allow users to customize the staging configuration, I propose providing them an in-app cartesian teleoperation interface to move the robot to a desired configuration. I also propose transparently exposing the technical constraint for this configuration to the user, by showing them the camera feed overlaid with the results of face detection. That way, the user can move their head through the range of poses they expect to be in during feeding and verify that face detection works from that staging configuration. Finally, I propose having a "Try It" button that enables users to experience a motion between that staging configuration and their mouth. Designed in this way, I believe users will be able to intuitively identify the staging configurations that aligns with their preferences and position relative to the robot.

Although not the focus of the user study, customizing how close the robot gets to the user's mouth is nonetheless necessary to allow users to fully customize the bite transfer. Thus, while customizing the staging configuration users will also be able to specify the distance to their mouth, for example with buttons or a text input to increase accessibility [26].

5.3.2 Robot-Driven Customization

Overview: I propose developing a robot-driven customization approach based on active learning from human feedback, a popular technique in HRI for learning models of user preferences [20; 106; 36]. Specifically, the robot will move to a staging configuration, optionally allow users to try the motion between the staging location and their mouth, and then ask users whether they like it or not. The robot will use that binary feedback to update its belief over which staging configurations the user prefers, and use that to inform the next staging configuration it tries.

Featurization: As mentioned in Sec. 5.1, a chief challenge in robot-driven customization is modeling the features Φ that influence user preferences. To do so, I ran an online pilot study that showed participants a variety of configurations the robot arm might be in and asked them qualitatively describe the robot and how it made them feel. $n = 11$ users participated, and the results revealed key features that impact users' preferences: whether the robot is centered on their mouth, how much of their visual field the robot occupies, how high the robot arm is relative to their face, whether the arm's bend in anthropomorphic, and more. I propose quantifying each of these (as in [22]) and using that to compute the features Φ of the proposed staging configuration.

User Preferences Model: I propose using a linear model of user preferences $h(f_\theta|w) = w \cdot \Phi(f_\theta)$, as this is commonly used in HRI [11; 92]. Then, the probability of a user saying that they like a staging configuration is $P(y_t = 1) = \sigma(h(f_{\theta_t}))$ where σ is the standard logistic function.

Belief Over User Preferences: I propose maintaining a multi-variate Gaussian belief $\mathcal{N}(w|\mu, \Sigma)$ over the user's preference vector w. Because we model feedback as the logistic of user preferences, there is no analytic solution for the posterior, so I propose using Laplace Approximations [72].

Learning Algorithm: I propose using Thompson Sampling [113], since it allows us to start with an informed prior and personalize it per user. Depending on the complexity of the feature space, we may not be able to analytically compute the staging configuration that optimizes the current guess of user preferences. Thus, I propose sampling staging configurations from which the user's face is visible and determining which of those optimizes the guess of user preferences.

5.4 Proposed User Study

Overview: I propose a user study comparing user-driven versus robot-driven customization of the robot's staging configuration. Specifically, participants with motor impairments will use both the robot-driven and user-driven customization systems, in a counter-balanced fashion.

Metrics: For each, I'll measure how long it takes them to reach a sufficiently-customized configuration. As in other studies on customization [129], I'll also ask users to rate their subjective experience such as how customized they feel the system is. Additionally, at multiple points while they are customizing the robot, I'll measuring their cognitive workload using the NASA-TLX [50]. Finally, I'll gather qualitative insights about users' perceived tradeoffs between the two through a semi-structured interview.

Hypotheses: I hypothesize that user-driven customization will result in a shorter time to customize the robot's behavior, and higher perceptions that the robot's behavior is customized to them. Further, robot-driven customization will result in lower overall workload but higher levels of frustration on the NASA-TLX, since users may get frustrated at how the process of mapping their feedback to parameter updates is opaque to them.

6
System Development & Evaluation (RQ4)

This chapter presents the 4th research foci: designing a deployable robot-assisted feeding system.

RQ4 How can we take a functional robot-assisted feeding system and make it deployable?

This research question was motivated by the observation that the space of personal physically assistive robots for people with disabilities has few in-the-wild deployments, with notable exceptions being [46; 64; 93; 61]. This is despite the increasing awareness in HRI that human perceptions of and interactions with robots in-the-wild often differs from those in-the-lab [62]. Part of the reason we see few deployments is that deployments introduce a multitude of off-nominal scenarios—scenarios where the user, robot, or environment behave differently from the system developers' expectations. Making a robot system that can work reliably in these many off-nominal scenarios is challenging, presenting a barrier deployability. Thus, I believe there is space for a research contribution focused on the key system design decisions that enable one to take a functional system and make it deployable, which I focus through robot-assisted feeding.

Nominal Scenario: a scenario where everything involved in system execution, including the user, robot, and environment, proceeds "according to plan" [83].

Off-Nominal Scenario: a scenario where something involved in system execution, including the user, robot, and environment, does not proceed according to plan [39].

In this chapter, I present key insights for taking our robot-assisted feeding system from a functional to a deployable state (ongoing work), and evaluations to assess system deployability.

6.1 *Insight #1: Users Can Resolve Off-Nominals, Given Control & Transparency*

In order to design a system that is robust to off-nominal scenarios, we must first understand those off-nominal scenarios. Thus, we worked with the community researcher who led the RQ1 interviews to enumerate off-nominals that can arise during robot-assisted feeding. This resulted in over 50 off-nominal scenarios, some of which are shown in in Table 6.1. These include situations where the user wants the robot to deviate from its typical behavior (e.g., the user is about to sneeze so wants the robot to wait), situations where the robot fails to do an autonomous behaviors (e.g., detect the user's face), and situations where the real environment does not align with the robot's assumption (e.g., the user and robot are angled with respect to the plate).

Recall from RQ1 that users want control over their robot-assisted feeding system: e.g., *"If I'm ready to eat and then someone starts talking to me, I'd want to [have the robot] wait until that person*

Table 6.1: A non-exhaustive list of off-nominal scenarios that can arise during robot-assisted feeding, and whether they are primarily caused by the user, robot, or environment.

User	Robot	Environment
User no longer wants the bite	Robot collides with object	Food falls off of the fork
User cannot eat (e.g., is coughing)	Robot fails to perceive bite	Plate moves (e.g., caregiver serves food)
User takes a partial bite	Robot fails to acquire bite	Local area network fails
User clicks unintended button	Robot stops far from face	Device running the web app fails
...	...	...

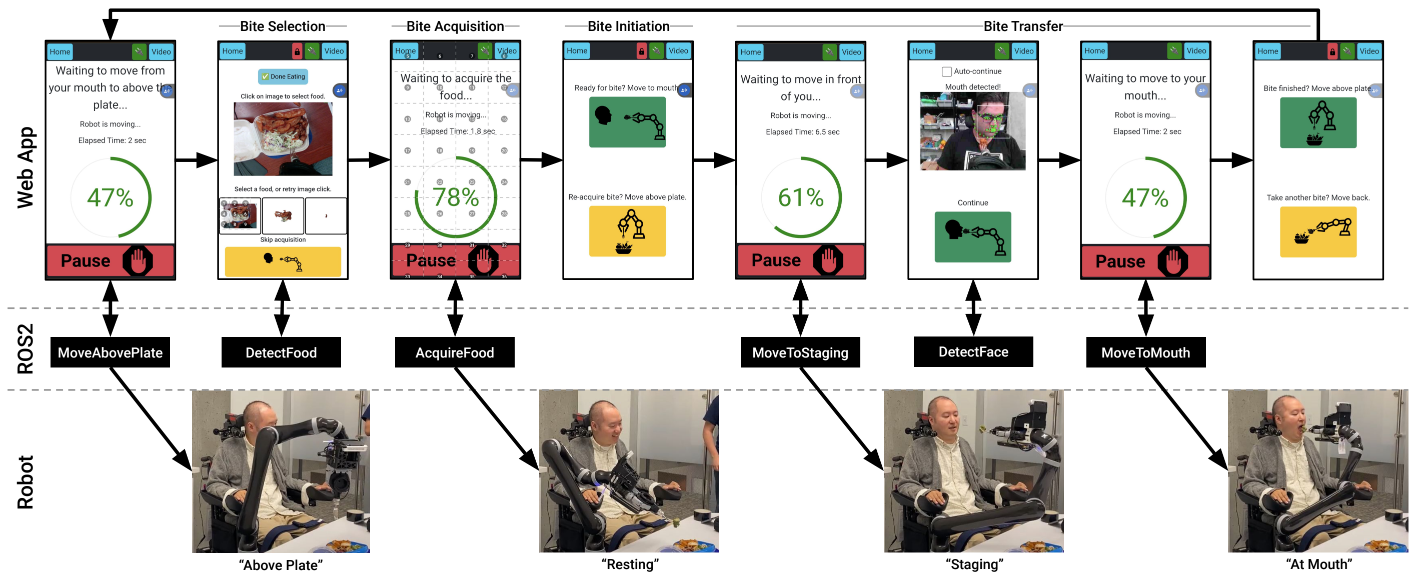

Figure 6.1: An overview of the robot-assisted feeding system's nominal state machine. Note that this doesn't show off-nominal states e.g., when the user preempts an action.

finishes" (P1). Users were articulating a desire to have control of their robot-assisted feeding system because they *expected off-nominal scenarios to arise*, and wanted the ability to *navigate the robot through those off-nominals* in the way they wanted. We draw upon this insight, by focusing on empowering users to navigate the robot through off-nominal scenarios back to nominal functioning.

6.1.1 A Software Architecture That Puts Users in Control

Providing high levels of user control starts with the software architecture. We architected our system so that the web app, not the robot, is in charge of system execution. Specifically, the robot code consists of modular ROS2 actions[1], such as "MoveAbovePlate" and "DetectFood," which the web app invokes. Each robot action only executes the specified behavior, and then returns to the web app. And then the user, through the web app, decides what action to invoke next. This architecture, shown in Fig. 6.1 ensures that no robot motion action will be executed without the user's explicit invokation, and that the user can preempt execution of an the action at any time.

[1] An action takes in a goal from the client, executes that goal asynchronously while sending feedback back to the client, and then returns a result. It also allows the client to preempt the action at any time.

6.1.2 A Web App That Provides Control & Transparency

Supervisory Control: The software architecture provides users the supervisory control [12] to stop robot actions, but what they want to do afterwards can vary. For example, say the user preempts the robot while it is approaching their mouth. From here, they may want it to resume (e.g., if they were about to sneeze), move to the resting configuration (e.g., if they got pulled into a conversation), or move above the plate (e.g., if the food fell off). The system cannot autonomously detect what the user wants it to do, so we must provide the user the option to decide (Fig. 6.2). In other words, although the system has a nominal flow through its state machine (Fig. 6.1), *providing users high flexibility to break that flow allows them to respond to a variety of off-nominal scenarios.*

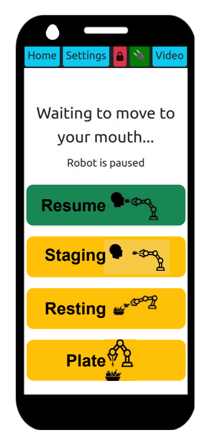

Figure 6.2: A proposed post-preemption screen.

Transparency: Providing users supervisory control allows them to address off-nominal scenarios that *they can detect*. However, empowering users to address off-nominals due to technical components of the system requires *providing them transparency into those system components*. For example, if the robot is approaching an object it might collide with, the user needs an transparency into how soon the robot will stop in order to decide whether to preempt it. We provide that transparency through a progress bar showing remaining robot motion (Fig. 6.1). As another example, if the F/T sensor loses WiFi connection, the user needs to know that so they can reboot it.

Ongoing work is focusing on rendering the status of all system components, hardware and software, to the user, and providing them controls to restart those system components. Importantly, *transperancy and control are two sides of the same coin* for empowering users to resolve off-nominals: they need the transparency to understand the problem and the control to address it.

Fallback Manual Control: Although the above cases give users supervisory control, all robot perception, planning, and controls is still handled by robot actions. However, if those actions repeatedly fail, providing users the option of fallback manual control can help. For example, if the robot repeatedly stops too far from the user's mouth, they can use an in-app teleoperation interface to move it the final way to their mouth. Ongoing work is focused on developing that teleoperation interface. Manual control can also apply in perception. For example, if food detection repeatedly fails to segment the user's preferred bite, they could draw a bounding box around it and use that to invoke bite acquisition. Thus, *fallback manual control should supplement robot actions to ensure users are always able to proceed, even in unlikely off-nominals.*

Caveats to High User Control: Although user control can get us far in terms of making the system robust to off-nominals, there are two key caveats. First, *users must be willing and able to exercise the control*. For example, conversations with participants who use the Kinova arm have revealed that teleoperating it to pick up a single bite of food can takes an unreasonable 5 to 40 minutes [18; 89]. Thus, if acquisition repeatedly fails, users are unlikely to use manual control to recover and are more likely to not use the system, necessitating a robust acquisition system (RQ2). All the above proposals for user control are ones that participants and community researchersexpressed willingeness to do. Second, *some users may be frustrated by too much control*. For example, while one user we spoke to wanted the robot to wait at the staging configuration until they tell it to move to their mouth, another wanted it to immediately move to their mouth to speed up the process. To account for these diverse preferences, I propose adding settings that allow users to auto-continue (e.g., Fig. 6.3), perhaps after a configureable number of seconds.

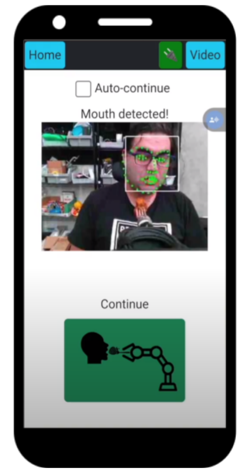

Figure 6.3: Users who don't want the robot to wait for them before moving to their mouth can check "Auto-continue."

6.1.3 *Flexible Robot Actions*

Providing users with high control necessitates flexible robot actions, because we lose guarantees over the robot state at the end of the action. For example, our original implementation of "MoveToStaging" required the fork to be flat, so food doesn't fall off. However, users might pause "AcquireFood" before the fork reaches a flat pose, and might subsequently invoke "MoveToStaging," which would immediately fail. Thus, *to accommodate high user control, robot actions must be designed to have fallbacks in case their preconditions are not met*. In the aforementioned case, we modified the robot action so if the fork is not flat, it removes the flatness constraint on the motion.

Similarly, the fact that users may use the system in any lighting condition, with any food type, necessitates an extremely flexible food perception system. Whereas past work used a food perception system that was trained on—and could mainly detect—6 food items [40], in the process of making the system deployable we substituted that with SegmentAnything [67], a flexible food segmentation system that has high success rates on diverse images. Thus, *developing a deployable system requires generalizable perception systems that are robust to environment and object variations*.

6.2 *Insight #2: Safety in All Levels of the System*

Safety is a crucial component of a deployable system, particularly one like the robot-assisted feeding system that operates extremely close to the user. While past robot-assisted feeding work had a researcher prepared to stop the robot as a fallback safety level [18], that is not an option for a deployable robot. Thus, we incorporate *safety into every level of our system*.

Web App: As mentioned above, at the app level users have the ability to preempt any robot action. The action then percolates the preemption into the code it is running. Since all robot motion actions are implemented as behavior trees [33], each behavior must already implement termination logic. Further, the modularity of behaviors makes it easy to test and verify termination logic. On visual inspection the robot stops immediately after the user clicks pause.

Motion Planning: On the motion planning level, we populate a dynamic collision map [56] with perceived depth camera obstacles. This ensures that all planned motion will avoid user- or environment-specific obstacles, e.g., assistive technology near the user or a cup on the table.

Controls: On the controls level, all controllers are force-gated. In other words, they subscribe to the force-torque sensor and stop motion as soon as a sensor reading exceeds a threshold. We keep these thresholds at conservative values during bite transfer—1N—and at values ranging from 4-35N during bite acquisition (these thresholds were learnt based on human data in RQ2).

Watchdog: As an overarching level of safety across all robot code, we use a watchdog architecture [29] to monitor the crucial safety components: the force-torque sensor and the physically emergency stop (e-stop) button mounted in a reachable location near the user's head. If the force-torque sensor stops publishing for at least 0.5s, or if the e-stop button is pressed, the watchdog "trips." Further, if a watchdog listener hasn't receive a message, or heartbeat, from the watchdog within 0.1s, they consider the watchdog to have tripped. All ROS2 actions have a watchdog listener and immediately abort if it trips, following the same termination logic as preemptions. Further, the process running the controllers also has a watchdog listener, and terminates itself if the watchdog "trips," thereby stopping all movement communication to the robot.

These layers of safety ensure the system is safe without researcher intervention. The user has two ways to immediately stop robot motion—through the app and with a physical e-stop button. Even if the user doesn't stop the robot, motion commands will stop being sent to the robot if: (a) a force-torque threshold is exceeded; or (b) the force-torque sensor stops sending readings.

6.3 *Insight #3: Portability is Key*

The robot-assisted feeding system must not hinder the user's ability to move their wheelchair, which means the system cannot require any wires leaving the wheelchair (e.g., to a power outlet) or an external internet connection (which may not be available everywhere). Thus, the system has its own local router and the router, laptop, and robot all draw power from the wheelchair's internal battery. Further, none of the sensors on the robot's arm have wires that leave the arm. The force-torque sensor is wireless, and the RGB-D camera connects to an Nvidia Jetson Nano computer mounted on the robot arm itself. Thus, the robot still has its full range of motion without concerns of wires getting caught and jerking the arm or sensors.

6.4 *Insight #4: Test in a Structured Progression*

To verify the system's deployability and catch any issues early, I propose testing it in a structured progression. In terms of test subjects, I propose starting with researchers without motor impairments, then researchers with motor impairments (the community researcher), before evaluating it with participants with motor impairments. In terms of environment, I propose moving from more structured to less structured: a conference room, then an atrium or cafeteria, then a user's home. In terms of food, I propose moving from easier to acquire foods like potato wedges to more difficult to acquire foods like cole slaw, noodles, and rice. In terms of the user's attentiveness during the meal, I propose moving from single-tasking (eating the meal) to solo multitasking (e.g., eating the meal while watching TV) to group multi-tasking (e.g., eating a meal while socializing). This structured testing progression will enable us to be confident in the system's functioning in earlier test situations before moving to more difficult ones.

6.5 *Evaluations*

So far, we have testing the system with a researcher with disabilities, in a structured out-of-lab environment (a conference room), with easier-to-acquire foods (Fig. 6.4). Once we finish thoroughly testing the system on the harder-to-acquire foods, I propose evaluating the system through multiple single-meal deployments and at least one week-long deployment.

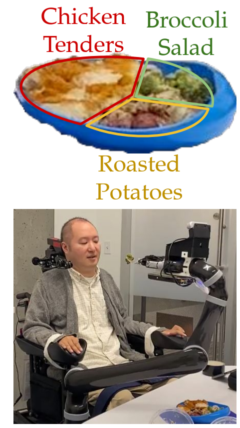

Figure 6.4: In a pilot study, the robot fed the community researcher a meal of chicken tenders, broccoli salad, and roast potatoes in a campus conference room.

6.5.1 *Multiple Single-Meal Deployments*

I propose running a system evaluation user study where $n \geq 5$ participants are fed by the robot-assisted feeding system in various campus environments outside of the lab (e.g., a conference room, atrium, cafeteria, etc.). After giving informed consent, the users will be introduced to the system and able to use the customization interface from RQ3 to tailor it to their needs. We will have already purchased their desired meal, and they will use the system to eat as much of it as desired. As they eat, we will track the number of times and reasons why we had to intervene (ideally 0), and periodically ask them to complete the NASA-TLX [50] to gauge their cognitive workload. At the end, users will rate their experience through the System Usability Scale [75] and we will conduct a semi-structured interview to understand more detailed feedback. This evaluation will help us measure our progress towards a deployable system, identify future work, and compare our system to others (through the use of validated scales).

6.5.2 *One Week-Long, In-Home Deployment*

I also intend to complete at least one home deployment, where the robot-assisted feeding system feeds one user, in their home, over an extended period of time (e.g., a week). When designing this study, I propose leveraging methodologies on $n = 1$ study design from personalized healthcare and behavioral science [123]. The core idea of $n = 1$ methodologies is for the user to serve as their own control. The results derive statistical power from many trials with one user, as opposed to one trial each with many users. For example, we could have the user alternate between meals fed by our system and by a caregiver, and measure metrics like their time to complete the meal, stress levels during the meal, feelings of satisfaction and self-efficacy, and how much caregiver time was required. By analyzing these metrics across the several meals the user eats over the course of the deployment, we can derive statistical results about how well the robot-assisted feeding system worked *for that user*, which is important because the ultimate goal of personal robotics is to develop robots that work well *for the specific people who are using it*. Further, note that because the home deployment will last an extended period of time, it is likely that some meals may need to be fed in-bed. If so, this will allow us to evaluate the system's ability to customize to the relative positions of the user, plate, and robot.

6.6 *Discussion: Extensibility Across Labs*

We are working with our collaborators at Cornell University to put our software architecture (Sec. 6.1.1) on their robot feeding system's hardware. Once it is done, the modularity of the architecture should enable smooth sharing of research advancements across the labs. For example, they developed an extension to "MoveToMouth" that goes inside the user's mouth (e.g., for users with no neck mobility); porting that to our system should just involve copying the action into our codebase. Similarly, to the best of my knowledge our food detection pipeline can detect a larger variety of foods than theirs, and porting that to their system will involve copying the action over. I look forward to this software architecture accelerating code-sharing in robot feeding research.

7
Conclusion

In this work, I proposed a research agenda that seeks to answer the following research question:

RQ-Thesis *How can we develop a deployable robot-assisted feeding system that can feed any user, in any environment, a meal of their choice, while aligning with their preferences?*

Specifically, I presented completed work that investigated: the needs and priorities of people with motor impairments when it comes to robot-assisted feeding (RQ1); and how to enable a robot-assisted feeding system to acquire the variety of foods users may want to eat (RQ2). I then presented proposed ongoing work on: ensuring the system can be customized to each user and environment (RQ3); and designing, developing, and evaluating a deployable system (RQ4).

Table 7.1 contains the proposed timeline for this thesis. I propose targeting HRI 2025 for RQ3, RSS 2023 for the single-meal deployments of RQ4, and perhaps a journal extension of that paper for the in-home deployment of RQ4.

7.1 Future Work

The proposed work will take us closer to widespread robot-assisted feeding being a reality for people with motor impairments, and paves the way for interesting future work.

Within bite acquisition, a key direction of future work is allowing users to specify their desired bite size. This is crucial to prevent choking, and is a make-or-break criteria for some users [89]. Within bite transfer, a key direction of future work is not obstructing the target of the user's gaze, such as a TV or their social dining companion. This may be possible with an in-app customization experience like those proposed in Ch. 5. Within system design, a key direction of future work is allowing users to seamelessly switch between autonomous feeding mode and Kinova's teleoperation mode. This will allow users to teleoperate the arm to get and heat food from the fridge, then turn on feeding mode and eat it, and then teleoperate the arm to put the dish in the sink.

Finally, achieving widespread adoption of a novel technology like robot-assisted feeding systems involves getting regulatory approval for semi-autonomous robots as assistive devices, securing health insurance coverage for robot-assisted feeding systems, and reducing the cost of these systems [35; 5; 111]. An alternative is releasing this technology as open-source do-it-yourself (DIY) kits [57], although this brings up questions of who has access to build such kits.

Table 7.1: Proposed timeline for this thesis.

Research Question(s)	End Quarter	Milestone
Needs Assessment (RQ1) & Acquisition (RQ2)	Autumn 2023	Pilot Single-Meal Deployment: 1
Customizability (RQ3) & Deployability (RQ4)	Winter 2024	Single-Meal Deployments: 5
Deployability (RQ4)	Spring 2024	In-home Deployment & Bed-side Feeding
Potential Internship	*Summer 2024*	*N/A*
Customizability (RQ3)	Autumn 2024	RQ3 Study
RQ-Thesis	Winter 2025	Dissertation & Defense

8 Bibliography

[1] Beeson automaddak feeder, Feb 2004. URL https://larimer.co.networkofcare.org/aging/assistive/detail.aspx?id=18902&cid=797&cn=Feeders&org=.

[2] Assistive dining technology 1970 through 2005. *Mealtime Partners, Inc.*, page 1–37, 2005.

[3] The mealtime partner dining system description, 2020. URL https://www.mealtimepartners.com/dining/mealtime-partner-dining-device.htm.

[4] Meal buddy, 2022. URL https://www.performancehealth.com/meal-buddy-systems.

[5] Gabriel Aguiar Noury, Andreas Walmsley, Ray B Jones, and Swen E Gaudl. The barriers of the assistive robotics market—what inhibits health innovation? *Sensors*, 21(9):3111, 2021.

[6] Patrícia Alves-Oliveira, Amal Nanavati, Tyler Schrenk, Ethan Gordon, Maya Cakmak, and Siddhartha S Srinivasa. Design principles for robot-assisted feeding in social contexts (supplementary materials), Dec 2022. URL osf.io/392hp.

[7] Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. Power to the people: The role of humans in interactive machine learning. *Ai Magazine*, 35(4):105–120, 2014.

[8] Paola Ardón, Maria E Cabrera, Eric Pairet, Ronald PA Petrick, Subramanian Ramamoorthy, Katrin S Lohan, and Maya Cakmak. Affordance-aware handovers with human arm mobility constraints. *IEEE Robotics and Automation Letters*, 6(2):3136–3143, 2021.

[9] Dar Artman, Mary Parkinson, and Pat Adanti-Joy. New obi robotic dining device a breakthrough for people living with physical challenges, Jul 2016. URL https://www.multivu.com/players/English/7850651-obi-robotic-dining-device/.

[10] Tommaso Lisini Baldi, Giovanni Spagnoletti, Mihai Dragusanu, and Domenico Prattichizzo. Design of a wearable interface for lightweight robotic arm for people with mobility impairments. In *International Conference on Rehabilitation Robotics (ICORR)*, pages 1567–1573. IEEE, 2017.

[11] Chandrayee Basu, Mukesh Singhal, and Anca D Dragan. Learning from richer human guidance: Augmenting comparison-based learning with feature queries. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 132–140, 2018.

[12] Jenay M. Beer, Arthur D. Fisk, and Wendy A. Rogers. Toward a framework for levels of robot autonomy in human-robot interaction. *J. Hum.-Robot Interact.*, 3(2):74–99, jul 2014. DOI: 10.5898/JHRI.3.2.Beer. URL https://doi-org.offcampus.lib.washington.edu/10.5898/JHRI.3.2.Beer.

[13] S. Belkhale, E.K. Gordon, Y. Chen, S. S. Srinivasa, T. Bhattacharjee, and D. Sadigh. Balancing efficiency and comfort in robot-assisted bite transfer. In *IEEE International Conference on Robotics and Automation*, 2022.

[14] Suneel Belkhale, Ethan K Gordon, Yuxiao Chen, Siddhartha Srinivasa, Tapomayukh Bhattacharjee, and Dorsa Sadigh. Balancing efficiency and comfort in robot-assisted bite transfer. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 4757–4763. IEEE, 2022.

[15] Jeanette Bell and Tuck Wah 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*, pages 1–13, 2019.

[16] Tapomayukh Bhattacharjee, Maria E Cabrera, Anat Caspi, Maya Cakmak, and Siddhartha S Srinivasa. A community-centered design framework for robot-assisted feeding systems. In *The 21st international ACM SIGACCESS conference on computers and accessibility*, pages 482–494, 2019.

[17] Tapomayukh Bhattacharjee, Gilwoo Lee, Hanjun Song, and Siddhartha S Srinivasa. Towards Robotic Feeding: Role of Haptics in Fork-based Food Manipulation. *IEEE Robotics and Automation Letters*, 2019.

[18] Tapomayukh Bhattacharjee, Ethan K Gordon, Rosario Scalise, Maria E Cabrera, Anat Caspi, Maya Cakmak, and Siddhartha 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)*, pages 181–190. IEEE, 2020.

[19] Steven W Brose, Douglas J Weber, Ben A Salatin, Garret G Grindle, Hongwu Wang, Juan J Vazquez, and Rory A Cooper. The role of assistive robotics in the lives of persons with disability. *American Journal of Physical Medicine & Rehabilitation*, 89(6):509–521, 2010.

[20] Maya Cakmak and Andrea L Thomaz. Designing robot learners that ask good questions. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, pages 17–24, 2012.

[21] Maya Cakmak, Crystal Chao, and Andrea L Thomaz. Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development*, 2(2):108–118, 2010.

[22] Maya Cakmak, Siddhartha S Srinivasa, Min Kyung Lee, Jodi Forlizzi, and Sara Kiesler. Human preferences for robot-human hand-over configurations. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1986–1993. IEEE, 2011.

[23] Gerard Canal, Guillem Alenyà, and Carme Torras. Personalization framework for adaptive robotic feeding assistance. In *International conference on social robotics*, pages 22–31. Springer, 2016.

[24] Gerard Canal, Guillem Alenyà, and Carme Torras. A taxonomy of preferences for physically assistive robots. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 292–297. IEEE, 2017.

[25] Gerard Canal, Carme Torras, and Guillem Alenyà. Are preferences useful for better assistance? a physically assistive robotics user study. *ACM Transactions on Human-Robot Interaction (THRI)*, 10(4):1–19, 2021.

[26] Scott Carter, Amy Hurst, Jennifer Mankoff, and Jack Li. Dynamically adapting guis to diverse input devices. In *Proceedings of the 8th international ACM SIGACCESS conference on Computers and accessibility*, pages 63–70, 2006.

[27] Dorota Chapko, Pino Frumiento, Nalini Edwards, Lizzie Emeh, Donald Kennedy, David McNicholas, Michaela Overton, Mark Snead, Robyn Steward, Jenny 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*, pages 1–17, 2020.

[28] Kevin Charras and Michèle Frémontier. Sharing meals with institutionalized people with dementia: a natural experiment. *Journal of Gerontological Social Work*, 53(5):436–448, 2010.

[29] Xi Chen, Juejing Feng, Martin Hiller, and Vera Lauer. Application of software watchdog as a dependability software service for automotive safety relevant systems. In *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07)*, pages 618–624. IEEE, 2007.

[30] Yang Chen, Diego Paez-Granados, Hideki Kadone, and Kenji Suzuki. Control interface for hands-free navigation of standing mobility vehicles based on upper-body natural movements. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11322–11329. IEEE, 2020.

[31] Sachin Chitta, Eitan Marder-Eppstein, Wim Meeussen, Vijay Pradeep, Adolfo Rodríguez Tsouroukdissian, Jonathan Bohren, David Coleman, Bence Magyar, Gennaro Raiola, Mathias Lüdtke, et al. ros_control: A generic and simple control framework for ros. *The journal of open source software*, 2(20):456–456, 2017.

[32] Cheng-Shiu Chung, Hongwu Wang, and Rory A Cooper. Functional assessment and performance evaluation for assistive robotic manipulators: Literature review. *The journal of spinal cord medicine*, 36(4):273–289, 2013.

[33] Michele Colledanchise and Petter Ögren. *Behavior trees in robotics and AI: An introduction*. CRC Press, 2018.

[34] Peter F Edemekong, Deb Bomgaars, Sukesh Sukumaran, and Shoshana B Levy. Activities of daily living. 2019.

[35] David Eguren and Jose Luis Contreras-Vidal. Navigating the fda medical device regulatory pathways for pediatric lower limb exoskeleton devices. *IEEE Systems Journal*, 15(2):2361–2368, 2020.

[36] Taylor A Kessler Faulkner, Elaine Schaertl Short, and Andrea L Thomaz. Interactive reinforcement learning with inaccurate feedback. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7498–7504. IEEE, 2020.

[37] Cara Bailey Fausset, Andrew J Kelly, Wendy A Rogers, and Arthur D Fisk. Challenges to aging in place: Understanding home maintenance difficulties. *Journal of Housing for the Elderly*, 25(2):125–141, 2011.

[38] Ryan Feng, Youngsun Kim, Gilwoo Lee, Ethan K Gordon, Matt Schmittle, Shivaum Kumar, Tapomayukh Bhattacharjee, and Siddhartha S Srinivasa. Robot-Assisted Feeding: Generalizing Skewering Strategies Across Food Items on a Plate. In *The International Symposium of Robotics Research*, pages 427–442. Springer, 2019.

[39] Donald Firesmith. The need to specify requirements for off-nominal behavior. *Software Engineering Institute Blog*, Jan 2012. URL https://insights.sei.cmu.edu/blog/the-need-to-specify-requirements-for-off-nominal-behavior/.

[40] Daniel Gallenberger, Tapomayukh Bhattacharjee, Youngsun Kim, and Siddhartha S Srinivasa. Transfer depends on acquisition: Analyzing manipulation strategies for robotic feeding. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 267–276. IEEE, 2019.

[41] Karen Glanz, Jessica J Metcalfe, Sara C Folta, Alison Brown, and Barbara Fiese. Diet and health benefits associated with in-home eating and sharing meals at home: A systematic review. *International journal of environmental research and public health*, 18(4):1577, 2021.

[42] E.K. Gordon*, A. Nanavati*, R. Challa, Bernie H. Zhu, Taylor A. Kessler Faulkner, and S. S. Srinivasa. Towards general single-utensil food acquisition with human-informed actions. In *Conference on Robot Learning*, 2023.

[43] Ethan K Gordon, Xiang Meng, Tapomayukh Bhattacharjee, Matt Barnes, and Siddhartha 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)*, pages 9659–9666. IEEE, 2020.

[44] Ethan K Gordon, Sumegh Roychowdhury, Tapomayukh Bhattacharjee, Kevin Jamieson, and Siddhartha 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)*, pages 10528–10535. IEEE, 2021.

[45] Jennifer Grannen, Yilin Wu, Suneel Belkhale, and Dorsa Sadigh. Learning bimanual scooping policies for food acquisition. *arXiv preprint arXiv:2211.14652*, 2022.

[46] Phillip M Grice and Charles C Kemp. In-home and remote use of robotic body surrogates by people with profound motor deficits. *PloS one*, 14(3):e0212904, 2019.

[47] Fabien Grzeskowiak, Louise Devigne, François Pasteau, Gustavo Souza Vieira Dutra, Marie Babel, and Sylvain Guégan. Swalkit: A generic augmented walker kit to provide haptic feedback navigation assistance to people with both visual and motor impairments. In *2022 International Conference on Rehabilitation Robotics (ICORR)*, pages 1–6. IEEE, 2022.

[48] Karen Hacker. *Community-based participatory research*. Sage publications, 2013.

[49] T Hammarberg. The right of people with disabilities to live independently and be included in the community. *Issue Paper, Council of Europe Commissioner for Human Rights*, 2012.

[50] Sandra G Hart. Nasa-task load index (nasa-tlx); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 50, pages 904–908. Sage publications Sage CA: Los Angeles, CA, 2006.

[51] Kelsey P Hawkins, Phillip M Grice, Tiffany L Chen, Chih-Hung King, and Charles C Kemp. Assistive mobile manipulation for self-care tasks around the head. In *2014 IEEE Symposium on computational intelligence in robotic rehabilitation and assistive technologies (CIR2AT)*, pages 16–25. IEEE, 2014.

[52] Laura V Herlant. Algorithms, implementation, and studies on eating with a shared control robot arm. 2018.

[53] Emmanuelle Hien and Bertrand L Deputte. Influence of a capuchin monkey companion on the social life of a person with quadriplegia: an experimental study. *Anthrozoös*, 10(2-3): 101–107, 1997.

[54] Michael Hillman. 2 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*, pages 25–44, 2004.

[55] Rachel Holladay, Shervin Javdani, Anca Dragan, and Siddhartha Srinivasa. Active comparison based learning incorporating user uncertainty and noise. In *RSS Workshop on Model Learning for Human-Robot Communication*, 2016.

[56] Armin Hornung, Kai M Wurm, Maren Bennewitz, Cyrill Stachniss, and Wolfram Burgard. Octomap: An efficient probabilistic 3d mapping framework based on octrees. *Autonomous robots*, 34:189–206, 2013.

[57] Amy Hurst and Jasmine Tobias. Empowering individuals with do-it-yourself assistive technology. In *The proceedings of the 13th international ACM SIGACCESS conference on Computers and accessibility*, pages 11–18, 2011.

[58] Sumio Ishii, Shinji Tanaka, and Fumiaki Hiramatsu. Meal assistance robot for severely handicapped people. In *Proceedings of 1995 IEEE International Conference on Robotics and Automation*, volume 2, pages 1308–1313. IEEE, 1995.

[59] Barbara A Israel, Amy J Schulz, Edith A Parker, and Adam B Becker. Review of community-based research: assessing partnership approaches to improve public health. *Annual review of public health*, 19(1):173–202, 1998.

[60] Shervin Javdani, Henny Admoni, Stefania Pellegrinelli, Siddhartha S Srinivasa, and J Andrew Bagnell. Shared autonomy via hindsight optimization for teleoperation and teaming. *The International Journal of Robotics Research*, 37(7):717–742, 2018.

[61] Rajat K Jenamani, ..., and Tapomayukh Bhattacharjee. Robot-assisted inside-mouth bite transfer using robust mouth perception and physical interaction-aware control. In *2024 19th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2024.

[62] Malte Jung and Pamela Hinds. Robots in the wild: A time for more robust theories of human-robot interaction, 2018.

[63] Sertac Karaman and Emilio Frazzoli. Incremental sampling-based algorithms for optimal motion planning. *Robotics Science and Systems VI*, 104(2):267–274, 2010.

[64] Seita Kayukawa, Daisuke Sato, Masayuki Murata, Tatsuya Ishihara, Hironobu Takagi, Shigeo Morishima, and Chieko Asakawa. Enhancing blind visitor's autonomy in a science museum using an autonomous navigation robot. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2023.

[65] Hyun K Kim, Heejin Jeong, Jangwoon Park, Jaehyun Park, Won-Seok Kim, Nahyeong Kim, Subin Park, and Nam-Jong Paik. Development of a comprehensive design guideline to evaluate the user experiences of meal-assistance robots considering human-machine social interactions. *International Journal of Human–Computer Interaction*, pages 1–14, 2022.

[66] JS Kim, W-K Song, SR Han, JE Heo, and KJ Lee. A study on the scooping and serving korean food using meal assistance robot. *Gerontechnology*, 21, 2022.

[67] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023.

[68] Klaus Krippendorff. *The content analysis reader*. Sage, 2009.

[69] Arun Kumar and Margaret Frances Phillips. Use of powered mobile arm supports by people with neuromuscular conditions. *Journal of Rehabilitation Research & Development*, 50 (1), 2013.

[70] Vijay Kumar, Tariq Rahman, and Venkat Krovi. Assistive devices for people with motor disabilities. *Wiley Encyclopedia of Electrical and Electronics Engineering*, 1997.

[71] Raja S Kushalnagar and Christian 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*, pages 1–6, 2020.

[72] Branislav Kveton, Manzil Zaheer, Csaba Szepesvari, Lihong Li, Mohammad Ghavamzadeh, and Craig Boutilier. Randomized exploration in generalized linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 2066–2076. PMLR, 2020.

[73] David Lee. *The Origins of an Everyday Behavior: Why do People Share Meals?* PhD thesis, The University of Mississippi, 2014.

[74] Min Kyung Lee, Sara Kiesler, Jodi Forlizzi, and Paul Rybski. Ripple effects of an embedded social agent: a field study of a social robot in the workplace. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 695–704, 2012.

[75] James R Lewis. The system usability scale: past, present, and future. *International Journal of Human–Computer Interaction*, 34(7):577–590, 2018.

[76] Ann-Louise Lindborg and Maria Lindén. Development of an eating aid–from the user needs to a product. In *pHealth 2015*, pages 191–198. IOS Press, 2015.

[77] Sara Ljungblad. Applying "designerly framing" to understand assisted feeding as social aesthetic bodily experiences. *ACM Transactions on Human-Robot Interaction*, 12(2):1–23, 2023.

[78] J. Tevere MacFayden. Educated monkeys help the disabled to help themselves. *Smithsonian*, 17:125–132, 1986.

[79] Dennis Maciuszek, Johan Aberg, and Nahid Shahmehri. What help do older people need? constructing a functional design space of electronic assistive technology applications. In *Proceedings of the 7th international ACM SIGACCESS conference on Computers and accessibility*, pages 4–11, 2005.

[80] Rishabh Madan, Rajat Kumar Jenamani, Vy Thuy Nguyen, Ahmed Moustafa, Xuefeng Hu, Katherine Dimitropoulou, and Tapomayukh Bhattacharjee. Sparcs: Structuring physically assistive robotics for caregiving with stakeholders-in-the-loop. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 641–648. IEEE, 2022.

[81] Jennifer Mankoff, Gillian R Hayes, and Devva 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)*, pages 3–10. ACM, 2010.

[82] Herbert L Meiselman. Dimensions of the meal: Science, culture, business, art, 2000.

[83] Merriam-Webster. Nominal. In *Merriam-Webster.com dictionary*. 2023. URL https://www.merriam-webster.com/dictionary/nominal.

[84] J Michaelis. Mechanical methods of controlling ataxia. *Bailliere's Clinical Neurology*, 2(1): 121–139, 1993.

[85] Meredith Minkler and Nina Wallerstein. *Community-based participatory research for health: From process to outcomes*. John Wiley & Sons, 2011.

[86] Abolfazl Mohebbi. Human-robot interaction in rehabilitation and assistance: a review. *Current Robotics Reports*, 1:131–144, 2020.

[87] Maia Naftali and Leah Findlater. Accessibility in context: understanding the truly mobile experience of smartphone users with motor impairments. In *Proceedings of the 16th international ACM SIGACCESS conference on Computers & accessibility*, pages 209–216, 2014.

[88] Amal Nanavati, Ramya Challa, Ethan K. Gordon, and Siddhartha S. Srinivasa. A Dataset of Food Manipulation Strategies for Diverse Foods, 2022. URL https://doi.org/10.7910/DVN/C8SI1D.

[89] Amal Nanavati, Patricia Alves-Oliveira, Tyler Schrenk, Ethan K. Gordon, Maya Cakmak, and Siddhartha S. Srinivasa. Design principles for robot-assisted feeding in social contexts. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, HRI '23, page 24–33, New York, NY, USA, 2023. ACM. ISBN 9781450399647. DOI: 10.1145/3568162.3576988. URL https://doi.org/10.1145/3568162.3576988.

[90] Amal Nanavati, Vinitha Ranganeni, and Maya 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*, 7, 2023.

[91] Isira Naotunna, Chamika Janith Perera, Chameera Sandaruwan, RARC Gopura, and Thilina Dulantha Lalitharatne. Meal assistance robots: A review on current status, challenges and future directions. In *2015 IEEE/SICE international symposium on system integration (SII)*, pages 211–216. IEEE, 2015.

[92] Heramb Nemlekar, Neel Dhanaraj, Angelos Guan, Satyandra K Gupta, and Stefanos Nikolaidis. Transfer learning of human preferences for proactive robot assistance in assembly tasks. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pages 575–583, 2023.

[93] Vy Nguyen. Increasing independence with stretch: A mobile robot enabling functional performance in daily activities. 2021.

[94] Niels Christian Mossfeldt Nickelsen. Imagining and tinkering with assistive robotics in care for the disabled. *Paladyn, Journal of Behavioral Robotics*, 10(1):128–139, 2019.

[95] Yutaro Ohshima, Yuichi Kobayashi, Toru Kaneko, Atsushi Yamashita, and Hajime Asama. Meal support system with spoon using laser range finder and manipulator. In *2013 IEEE workshop on robot vision (WORV)*, pages 82–87. IEEE, 2013.

[96] Tomohiro Oka, Jorge Solis, Ann-Louise Lindborg, Daisuke Matsuura, Yusuke Sugahara, and Yukio Takeda. Kineto-elasto-static design of underactuated chopstick-type gripper mechanism for meal-assistance robot. *Robotics*, 9(3):50, 2020.

[97] Jan Ondras, Abrar Anwar, Tong Wu, Fanjun Bu, Malte Jung, Jorge Jose Ortiz, and Tapomayukh Bhattacharjee. Human-robot commensality: Bite timing prediction for robot-assisted feeding in groups. *arXiv preprint arXiv:2207.03348*, 2022.

[98] Juan F Orejuela-Zapata, Sarita Rodríguez, and Gonzalo Llano Ramírez. Self-help devices for quadriplegic population: a systematic literature review. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 27(4):692–701, 2019.

[99] Daehyung Park, Hokeun Kim, Yuuna Hoshi, Zackory Erickson, Ariel Kapusta, and Charles 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)*, pages 5406–5413. IEEE, 2017.

[100] Daehyung Park, Yuuna Hoshi, Harshal P Mahajan, Ho Keun Kim, Zackory Erickson, Wendy A Rogers, and Charles C Kemp. Active robot-assisted feeding with a general-purpose mobile manipulator: Design, evaluation, and lessons learned. *Robotics and Autonomous Systems*, 124:103344, 2020.

[101] Max Pascher, Annalies Baumeister, Stefan Schneegass, Barbara Klein, and Jens 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*, pages 331–351. Springer, 2021.

[102] Karie Jo Peralta. Politics of knowledge in community-based work. In *Dimensions of Community-Based Projects in Health Care*, pages 67–78. Springer, 2018.

[103] Chamika Janith Perera, Thilina Dulantha Lalitharatne, and Kazuo 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)*, pages 1760–1765. IEEE, 2017.

[104] Guy N Phillips. Feasibility study for assistive feeder. *Southwest Research Institute*, page 1–130, 1987.

[105] Camilo Perez Quintero, Oscar Ramirez, and Martin Jägersand. Vibi: Assistive vision-based interface for robot manipulation. In *IEEE International Conference on Robotics and Automation*, pages 4458–4463. IEEE, 2015.

[106] Mattia Racca, Antti Oulasvirta, and Ville Kyrki. Teacher-aware active robot learning. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 335–343. IEEE, 2019.

[107] Vinitha Ranganeni, Noah Ponto, and Maya Cakmak. Evaluating customization of remote tele-operation interfaces for assistive robots. *arXiv preprint arXiv:2304.02771*, 2023.

[108] Vinitha Ranganeni, Mike Sinclair, Eyal Ofek, Amos Miller, Jonathan Campbell, Andrey Kolobov, and Edward Cutrell. Exploring levels of control for a navigation assistant for blind travelers. *arXiv preprint arXiv:2301.02336*, 2023.

[109] Mitchel Resnick and Eric Rosenbaum. Designing for tinkerability. *Design, make, play: Growing the next generation of STEM innovators*, pages 163–181, 2013.

[110] Travers Rhodes and Manuela Veloso. Robot-driven trajectory improvement for feeding tasks. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2991–2996. IEEE, 2018.

[111] GRBE Romer, Harry JA Stuyt, and Albér Peters. Cost-savings and economic benefits due to the assistive robotic manipulator (arm). In *9th International Conference on Rehabilitation Robotics, 2005. ICORR 2005.*, pages 201–204. IEEE, 2005.

[112] Stephanie L. Rosenthal. *Human-Centered Planning for Effective Task Autonomy*. PhD thesis, Carnegie Mellon University, USA, 2012. URL http://www.cs.cmu.edu/~mmv/papers/RosenthalThesis.pdf.

[113] Daniel J Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, Zheng Wen, et al. A tutorial on thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018.

[114] 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*, 22(1):38–57, 1985.

[115] Teresa E Seeman. Social ties and health: The benefits of social integration. *Annals of epidemiology*, 6(5):442–451, 1996.

[116] Lorenzo Shaikewitz, Yilin Wu, Suneel Belkhale, Jennifer Grannen, Priya Sundaresan, and Dorsa Sadigh. In-mouth robotic bite transfer with visual and haptic sensing. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9885–9895. IEEE, 2023.

[117] Shoshanna Sofaer. Qualitative methods: what are they and why use them? *Health services research*, 34(5 Pt 2):1101, 1999.

[118] Won-Kyung Song and Jongbae Kim. Novel assistive robot for self-feeding. *Robotic Systems-Applications, Control and Programming*, pages 43–60, 2012.

[119] Won-Kyung Song, Won-Jin Song, Yale Kim, and Jongbae Kim. Usability test of knrc self-feeding robot. In *2013 IEEE 13th International Conference on Rehabilitation Robotics (ICORR)*, pages 1–5. IEEE, 2013.

[120] Ryoji Soyama, Sumio Ishii, and Azuma Fukase. The development of meal-assistance robot'my spoon'. In *Proceedings of the 8th International Conference on Rehabilitation Robotics, 2003*, pages 88–91, 2003.

[121] Priya Sundaresan, Suneel Belkhale, and Dorsa Sadigh. Learning visuo-haptic skewering strategies for robot-assisted feeding. In *Conference on Robot Learning*, pages 332–341. PMLR, 2023.

[122] Priya Sundaresan, Jiajun Wu, and Dorsa Sadigh. Learning sequential acquisition policies for robot-assisted feeding. *arXiv preprint arXiv:2309.05197*, 2023.

[123] Robyn L Tate and Michael Perdices. N-of-1 trials in the behavioral sciences. *The essential guide to N-of-1 trials in health*, pages 19–41, 2015.

[124] Noriyuki Tejima. Evaluation of rehabilitation robots for eating. In *Proceedings 5th IEEE International Workshop on Robot and Human Communication. RO-MAN'96 TSUKUBA*, pages 118–120. IEEE, 1996.

[125] Kristina A Theis, Amy Steinweg, Charles G Helmick, Elizabeth Courtney-Long, Julie A Bolen, and Robin Lee. Which one? what kind? how many? types, causes, and prevalence of disability among us adults. *Disability and health journal*, 12(3):411–421, 2019.

[126] Title and authors omitted for anonymous review. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2022.

[127] Mike Topping. An overview of the development of handy 1, a rehabilitation robot to assist the severely disabled. *Artificial Life and Robotics*, 4(4):188–192, 2000.

[128] Jennifer Utter, Simon Denny, Roshini Peiris-John, Emma Moselen, Ben Dyson, and Terryann Clark. Family meals and adolescent emotional well-being: findings from a national study. *Journal of nutrition education and behavior*, 49(1):67–72, 2017.

[129] Lex Van Velsen, Thea Van Der Geest, Rob Klaassen, and Michael Steehouder. User-centered evaluation of adaptive and adaptable systems: a literature review. *The knowledge engineering review*, 23(3):261–281, 2008.

[130] Meera Viswanathan, Alice Ammerman, Eugenia Eng, Gerald Garlehner, Kathleen N Lohr, Derek Griffith, Scott Rhodes, Carmen Samuel-Hodge, Siobhan Maty, Linda Lux, et al. Community-based participatory research: Assessing the evidence: Summary. *AHRQ evidence report summaries*, 2004.

[131] Jacob O Wobbrock, Shaun K Kane, Krzysztof Z Gajos, Susumu Harada, and Jon Froehlich. Ability-based design: Concept, principles and examples. *ACM Transactions on Accessible Computing (TACCESS)*, 3(3):1–27, 2011.

[132] *Disability*. World Health Organization, Mar 2023. URL https://www.who.int/news-room/fact-sheets/detail/disability-and-health.

[133] Akira Yamazaki and Ryosuke Masuda. Autonomous foods handling by chopsticks for meal assistant robot. In *ROBOTIK 2012; 7th German Conference on Robotics*, pages 1–6. VDE, 2012.

[134] Ruolin Ye, Wenqiang Xu, Haoyuan Fu, Rajat Kumar Jenamani, Vy Nguyen, Cewu Lu, Katherine Dimitropoulou, and Tapomayukh Bhattacharjee. Rcare world: A human-centric simulation world for caregiving robots. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 33–40. IEEE, 2022.

[135] Yan Zhang, Ziang Li, Haole Guo, Luyao Wang, Qihe Chen, Wenjie Jiang, Mingming Fan, Guyue Zhou, and Jiangtao 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*, pages 1–22. ACM, 2023.