Achieving Deployable Autonomy through Customizability and Human-in-the-Loop: A Case Study in Robot-assisted Feeding

Amal Nanavati
amaln@cs.washington.edu
University of Washington
Seattle, WA, USA

ABSTRACT

Despite decades of research on personal physically assistive robots for people with motor impairments, deployments of such robots are still few. Part of the reason is that every user's needs, environments, and care routines are unique, making it difficult to develop a sufficiently customized and robust robot. I present past and ongoing research with the ultimate aim of enabling a robot-assisted feeding system to feed a meal to any user, in any environment, without researcher intervention, in a way that aligns with the user's preferences. Our key insight is that the robot and user form a joint human-robot system that is working together to feed the user. Thus, we can achieve deployable autonomy by providing the user with intuitive and transparent controls to: customize the robot to their needs and environment; and make the robot's execution robust.

CCS CONCEPTS

• **Computer systems organization** → **Robotic autonomy**; • **Human-centered computing** → **Accessibility technologies**.

KEYWORDS

robot-assisted feeding, human-robot interaction, accessibility

ACM Reference Format:
Amal Nanavati. 2024. Achieving Deployable Autonomy through Customizability and Human-in-the-Loop: A Case Study in Robot-assisted Feeding. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24 Companion), March 11–14, 2024, Boulder, CO, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3610978.3638356

1 INTRODUCTION AND MOTIVATION

Research since the 1980s [24] has focused on developing personal physically assistive robots for people with motor impairments. Other than a few exceptions [8, 9, 13, 22], such robots tend not to be deployed outside the lab. Part of the reason is that every user's needs, environments, and care routines are unique, making it difficult to develop a sufficiently customized and robust robot.

In this work, I focus on robot-assisted feeding (RAF) as a case study for making physically assistive robots sufficiently customized and robust to work across diverse users and environments. Our goal is to enable an RAF system to feed a meal to any user, in any environment, without researcher intervention, in a way that aligns with their preferences. I investigate the following questions:

RQ1: *Users' Needs Assessment*: What features do participants desire from their RAF system? [Completed]

RQ2: *Customizability*: How can an RAF system customize to users' environments and preferences? [Ongoing]

RQ3: *Robustness*: How can an RAF system feed users without researcher intervention? [Ongoing]

Our key insight is that because the user is co-located with the robot, desires control over it, and has full alignment with its goals, we can view them as a joint human-robot system. Thus, we can achieve deployable autonomy by giving the user intuitive and transparent controls to: (a) customize the robot to their needs and environment; and (b) make the robot robust to off-nominal scenarios.

2 USERS' NEEDS ASSESSMENT (RQ1)

To investigate **RQ1**, we worked with a community researcher[1] to investigate: (a) the challenges people with motor impairments face during social dining; and (b) how we should design a robot-assisted feeding system to enable meaningful social dining[2] [17]. Specifically, we conducted interviews with 10 participants with diverse motor impairments. We discussed participants' current social dining practices, showed them videos of an RAF system being used in social settings (Fig 1a), and discussed on how participants' would like RAF systems to be designed. Fig 1b shows illustrative quotes for two design principles that emerged from this work.

A key insight is that **users seek personalization and consistency** in their feeding experience. This is motivated by their needs (e.g., small bites to prevent choking), preferences (e.g., wanting to be fed from their dominant side), and environments (e.g., wanting to be fed quickly at home and leisurely when socializing). Some caregivers provide personalization, but the reality of multiple caregivers precludes consistent personalization. This motivates **RQ2**.

Another key insight is that **users desire control over their RAF system**. This is partly due to a learnt resignation to technology errors (Fig 1b.C). It is also due to the expectation that situations will arise where the technology should deviate from nominal behavior, and users are best-suited to address that: *"If I'm ready to eat and then someone starts talking to me, I'd want to [have the robot] wait until that person finishes"* (P1). This insight reveals that human-in-the-loop control can be a way to address robustness in **RQ3**.

3 CUSTOMIZABILITY (RQ2)

Without customizing to users' needs and environments, an RAF system will not work. Consider the following customization realms:

A **the bite transfer**, e.g., some users can only chew from a particular side; some users can't rotate their necks.

B **the spatial arrangement of the user, robot, and plate**, e.g., some wheelchairs don't fit under tables so users have to sit sideways relative to the plate; some users need to be fed in bed.

HRI '24 Companion, March 11–14, 2024, Boulder, CO, USA

ACM ISBN 979-8-4007-0323-2/24/03
https://doi.org/10.1145/3610978.3638356

[1] In community-based participatory research (CBPR), academics researchers work equitably with community members throughout all research stages. CBPR is used in health [12, 26] and assistive technology [2, 6, 15, 17] research.

[2] Needs assessments that are not specific to social dining can be found in [3, 14, 19].

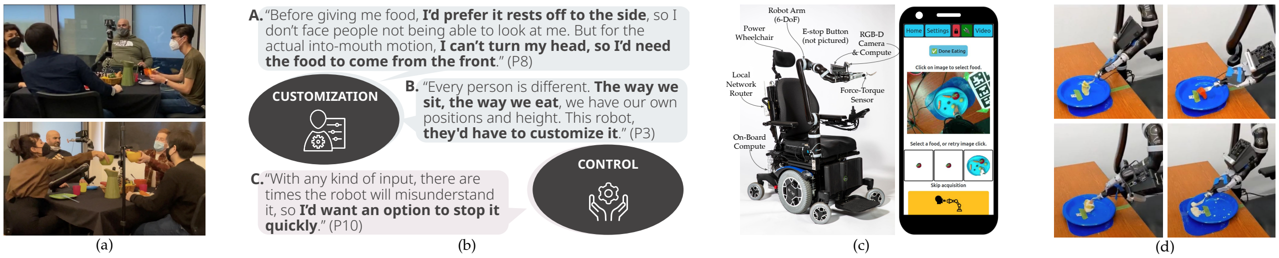

Figure 1: (a) The robot-assisted feeding (RAF) system being used during social dining [17]. (b) Users want to customize and control their RAF system [17]; (c) The RAF system consists of a wheelchair-mounted robot arm, a web app, and the user. (d) The author's past work enabled the system to acquire food like jello, sandwich bites, and mashed potatoes with $\geq 80\%$ accuracy [7].

One approach to customization has the robot collect data from users and learn parameters of its behavior that it thinks align with user preferences (i.e., robot-driven customization) [5]. Although this has shown promising results, robot-driven customization can leave users feeling disempowered or insufficiently included [1]. Our key insight is that users are experts at what they want; by providing intuitive knobs to directly modify robot parameters, we can empower them to customize their RAF system (i.e., user-driven customization). The key challenges to developing user-driven customization are: (1) identifying robot parameters that are expressive enough to capture user preferences, while not being unintuitive; (2) providing sufficient transparency into how parameters change robot behavior so users can make informed customization decisions.

As an initial investigation, we ran an online pilot study (n=11) to understand the features of robot arm motion that impact participants' bite transfer preferences. This revealed the importance of the robot not blocking the user's visual field, and the diversity of preferences—e.g., some wanted it to approach from the front whereas others preferred the side. To capture these preferences, I propose allowing users to customize the configuration the robot arm takes before transfer. Since the robot moves in a straight line to the mouth after that, customizing that configuration will allow users to reflect diverse preferences: approaching from a desired angle, avoiding their visual field, etc. I then propose running a user study comparing user-driven and robot-driven customization. The user-driven approach will be informed by "Designing for Tinkerability" [21] to ensure intuitive convergence to a desired configuration, and the robot-driven approach will be informed by active learning methods [5, 20]. My hypothesis is that user-driven customization will allow users to converge to their preferred robot behavior faster, and allow them to capture preferences that may be hard to model.

4 ROBUSTNESS (RQ3)

For users to use RAF systems without researcher intervention, the system must to be robust to the off-nominal scenarios that arise outside the lab. Thus, we worked with the community researcher to identify off-nominals that can arise during robot-assisted feeding. This resulted in over 50 off-nominals, including situations where the user wants the robot to deviate from its nominal behavior (e.g., they are about to sneeze and want the robot to wait), the robot fails to do an autonomous behavior (e.g., detect the user's face), or the real environment doesn't align with the robot's assumptions (e.g., the caregiver moves the plate while serving food). While some off-nominals can be addressed through autonomous recovery behaviors, this faces the catch-22 that those recovery behaviors can also fail. Instead, we propose designing a human-in-the-loop system that empowers *users* to navigate the robot through off-nominal scenarios to restore nominal functioning.

Insight #1: Human Control By Design Many human-in-the-loop systems are "robot control" by design, where the robot controls system execution and waits for human inputs at pre-defined times [4, 18, 23, 25]. In contrast, we develop a "human control" by design system, where the robot exposes an API of modular actions (e.g., "acquire food," "move to mouth") and the user interface—a web app (Fig 1c)—invokes each action when the user specifies. This provides the user wide latitude to decide which action to invoke, when to invoke it, and whether to preempt the executing action.

Insight #2: Autonomous Safety Checks User control is one part of making the system robust, but it takes users time to detect and respond to off-nominals. For the most safety-critical off-nominals—collision with the environment or user—we also run faster autonomous safety checks. These checks are implemented in all system layers: a dynamic collision map [11] in the planning layer to avoid unmodeled obstacles (e.g., assistive technology around the user's face); force-gated controllers in the execution layer that shut down upon unexpectedly large forces; and a software watchdog that terminates robot motion if any safety subsystem dies.

Insight #3: Transparency for Exercising Control Autonomous system components can fail; empowering users to resolve those requires transparency into system functionality and the ability to execute recovery behaviors. Consider when the user sees the robot approach an obstacle (e.g., a drink). Users need transparency into how soon the robot will stop to assess whether to preempt it. Once preempted, re-invoking the action may cause the same issue; instead, users need fallback teleoperation control to move the robot around the obstacle. Overall, users need the transparency to understand system problems and the control to address them.

Evaluation System development is mostly complete; the first two insights are implemented and the third is in-progress. I propose evaluating the complete RAF system with people with motor impairments in varied environments, e.g., a cafeteria, at home in front of a TV, sitting on their bed. One metric will be the number of researcher interventions. Other metrics will be time per bite, user's cognitive workload during eating [10], and user's system usability rating [16]. My hope is that with self-guided customization and in-app controls, users can have the robot feed them without researcher intervention in a way that aligns with their preferences.

ACKNOWLEDGMENTS

This paper is dedicated to the memory of Tyler Schrenk, a dear collaborator, advocate, and friend. This work was (partially) funded by: the NSF (DGE-1762114, #2132848, #2007011); the ONR (#N00014-17-1-2617-P00004 and #2022-016-01 UW); and Amazon.

REFERENCES

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

[2] Jeanette Bell and Tuck Wah Leong. 2019. 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*. 1–13.

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

[4] Tapomayukh Bhattacharjee, Ethan K Gordon, Rosario Scalise, Maria E Cabrera, Anat Caspi, Maya Cakmak, and Siddhartha S Srinivasa. 2020. Is more autonomy always better? exploring preferences of users with mobility impairments in robot-assisted feeding. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*. 181–190.

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

[6] Dorota Chapko, Pino Frumiento, Nalini Edwards, Lizzie Emeh, Donald Kennedy, David McNicholas, Michaela Overton, Mark Snead, Robyn Steward, Jenny M Sutton, et al. 2020. " 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*. 1–17.

[7] E.K. Gordon*, A. Nanavati*, R. Challa, Bernie H. Zhu, Taylor A. Kessler Faulkner, and S. S. Srinivasa. 2023. Towards General Single-Utensil Food Acquisition with Human-Informed Actions.

[8] Phillip M Grice and Charles C Kemp. 2016. Assistive mobile manipulation: Designing for operators with motor impairments. In *RSS 2016 Workshop on Socially and Physically Assistive Robotics for Humanity*.

[9] Fabien Grzeskowiak, Louise Devigne, François Pasteau, Gustavo Souza Vieira Dutra, Marie Babel, and Sylvain Guégan. 2022. 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)*. IEEE, 1–6.

[10] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*. Vol. 52. Elsevier, 139–183.

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

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

[13] Seita Kayukawa, Daisuke Sato, Masayuki Murata, Tatsuya Ishihara, Hironobu Takagi, Shigeo Morishima, and Chieko Asakawa. 2023. 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*. 1–14.

[14] Hyun K Kim, Heejin Jeong, Jangwoon Park, Jaehyun Park, Won-Seok Kim, Nahyeong Kim, Subin Park, and Nam-Jong Paik. 2022. 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* 38, 17 (2022), 1687–1700.

[15] Raja S Kushalnagar and Christian Vogler. 2020. Teleconference accessibility and guidelines for deaf and hard of hearing users. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*. 1–6.

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

[17] Amal Nanavati, Patricia Alves-Oliveira, Tyler Schrenk, Ethan K. Gordon, Maya Cakmak, and Siddhartha S. Srinivasa. 2023. Design Principles for Robot-Assisted Feeding in Social Contexts. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Stockholm, Sweden) *(HRI '23)*. ACM, New York, NY, USA, 24–33. https://doi.org/10.1145/3568162.3576988

[18] Amal Nanavati, Nick Walker, Lee Taber, Christoforos Mavrogiannis, Leila Takayama, Maya Cakmak, and Siddhartha Srinivasa. 2022. Not all who wander are lost: A localization-free system for in-the-wild mobile robot deployments. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 422–431.

[19] Max Pascher, Annalies Baumeister, Stefan Schneegass, Barbara Klein, and Jens Gerken. 2021. Recommendations for the Development of a Robotic Drinking and Eating Aid-An Ethnographic Study. In *Human-Computer Interaction–INTERACT 2021: 18th IFIP TC 13 International Conference, Bari, Italy, August 30–September 3, 2021, Proceedings, Part I 18*. Springer, 331–351.

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

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

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

[23] Stefanie Tellex, Ross Knepper, Adrian Li, Daniela Rus, and Nicholas Roy. 2014. Asking for help using inverse semantics. (2014).

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

[25] Manuela Veloso, Joydeep Biswas, Brian Coltin, and Stephanie Rosenthal. 2015. Cobots: Robust symbiotic autonomous mobile service robots. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.

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