

My research aims to enable robots to personalize their assistance to user needs. With aging populations and healthcare shortages, robots hold great potential to support older adults at home, by offering physical assistance for dressing and mobility, providing cognitive support through reminders, and encouraging social engagement. Beyond eldercare, robots can also benefit K-12 education, mental health support, and rehabilitation. Because each person's needs are unique and evolving, robots must continuously adjust their assistance. For example, a dressing robot might detect its user's limited shoulder flexibility and modify its movements to avoid causing discomfort.

To personalize robot assistance, my research focuses on interpreting the subtle cues that people naturally provide. Direct guidance, like manually showing a robot how to help, requires time and effort, making it impractical for long-term use. In contrast, people's everyday behaviors around robots reveal valuable insights into their needs, serving as "implicit human feedback". For example, a person's relaxed posture and continued engagement with a robot may indicate comfort. With advancements in movement tracking and wearable technologies, robots can now measure a wide range of such feedback. My research develops algorithms that enable robots to interpret these signals, understand user needs, and adjust assistance accordingly.

The challenge with implicit feedback is its ambiguity; the same behavior can mean different things. For instance, if a person moves more actively during dressing, they may be seeking more control over the dressing process, so the robot should give them space, or they may be in a hurry, so the robot should speed up. When interpreting implicit feedback, robots must assess their uncertainty, or how much they don't yet understand about the user's needs, and then either act cautiously to respect this uncertainty or adjust their actions to gain a better understanding.

To help robots interpret implicit feedback while managing uncertainty, my research takes an interdisciplinary approach, combining human decision-making models from psychology with algorithms from machine learning and control theory. For example, when a robot asks users multiple-choice questions to learn their preferences, such as whether it should start dressing the left or right arm, my recent work [1] was the first to use response times (how quickly users answer) as implicit feedback. Psychological research shows that faster responses often indicate stronger preferences. I developed a machine learning algorithm that uses response times to infer the strength of user preferences, enabling robots to learn user preferences more efficiently. Statistical analysis shows that this algorithm is especially effective when the robot recognizes that users have strong preferences, as this recognition allows the robot to trust its understanding and reduce its uncertainty about the user preferences, ultimately enhancing preference learning. This work bridges decades of psychological research on response times with preference learning methods from economics and statistics. Moving forward, I aim to strengthen the connections between human-centered studies and algorithmic development, fostering interdisciplinary collaboration.

I have demonstrated my work in various real-world applications, as shown in fig. 1. I led three projects [5, 7, 8] to become permanent exhibits at the MIT Museum, recognized for their technical innovation, engaging interactions, and educational impact. Looking ahead, I aim to further enhance human quality of life by developing robots that provide cognitive and physical assistance for diverse users, from daily assistance for older adults to personalized tutoring for children.




Figure 1: Real-world applications of my past research, including robot-assisted dressing [2, 3], collaborative meal preparation [4, 5], and industrial assembly [6].

RESEARCH CONTRIBUTIONS

My research enables robots to interpret human implicit feedback, translate it into ongoing **cognitive and behavioral processes**, extract critical information like preferences and intentions, and adapt their assistance accordingly. When robots seek to learn user preferences through multiple-choice questions, the underlying **cognitive process** reveals how strong or certain a user's preference is. However, these cognitive processes are not directly observed by the robot. My contribution #1 [1] was the first to use response times as implicit feedback,

alongside user choices, to uncover users' decision-making processes. This approach allowed robots to learn user preferences more efficiently with fewer interactions.

In physical assistance tasks, understanding the user's **behavioral process** is essential for safe and comfortable support. However, parts of the human body, like an elbow during dressing, are often obscured from the robot's view. My contribution #2 [3] was the first to enable robots to assess their uncertainty about human behavior and reliably estimate the positions of these obscured body parts using implicit feedback, such as force readings.

For multi-step tasks, such as collaborative meal preparation, robots need to anticipate the user's intended sequence of steps as part of the user's **cognitive process**. However, these evolving intended steps are not only unobserved by the robot but can also branch into multiple next steps. In contribution #3 [4], with collaborators, I was the first to enable robots to use implicit feedback, such as hand movements, to assess uncertainty about these branching intentions and adjust their actions to better understand user intentions and complete the task.

#1 Revealing User Cognitive Processes for Efficient Preference Learning. Understanding user preferences is key for effective robot assistance. To learn user preferences, a common framework is to ask users to choose between two options and collecting data over time. For example, a robotic wheelchair might ask users to choose between two routes to learn their navigation preferences. This framework is also widely used in recommender systems and large language models to align with user preferences. Current state-of-the-art methods learn preferences solely from users' choices, often modeled using the classic Bradley and Terry [9] from 1952. However, these choices only represent the outcomes of users' decision-making processes and provide limited information about the strength of preferences, which restricts the efficiency of preference learning.

My work [1] was the first to incorporate response times to reveal the cognitive process behind users' choices and leverage it for preference learning. Drawing on the Drift-Diffusion Model [10] from 1978, I developed an algorithm that uses choices (explicit feedback¹) to identify the user's preferred option and response times (implicit feedback) to assess the strength of those preferences. Statistical analysis shows that incorporating response times significantly enhances preference learning when the robot recognizes that users have strong preferences. This recognition, based on response times, allows the robot to trust its understanding of user preferences inferred from choices, reducing uncertainty and improving the learning process. Empirical results from three simulation studies demonstrate that my method, by incorporating response times, more accurately identified users' preferred options, reducing the chance of misidentification by up to 55 %. This work bridges over 50 years of psychological research on response times with preference learning methods, offering a novel approach that can potentially accelerate how robots, recommender systems, and large language models align with user preferences.

#2 Estimating User Behavior for Safe Assistance. In physical assistance tasks, parts of the human body are often obscured from the robot's view. Estimating these obscured body parts using sensor data is essential for safe assistance. For example, a robot can use force readings to estimate the position of an elbow during dressing. Set-based methods from control theory estimate the position of a body part as a point within a box that represents the robot's uncertainty about its estimation (see the green box in the figure on the right for illustration). These methods theoretically guarantee that the true body part lies within the box, allowing the robot to consider all possible positions and potentially plan safe assistance. However, these methods rely on perfect, manually specified models of human behavior. Since individual behavior varies, manually specifying these models for each person is impractical.




My work [3] introduces the first set-based method that learns human behavior models from pre-collected data while assessing the robot's uncertainty about these models due to scarce and noisy data. When interacting with a user in real time, my method dynamically adjusts the size of the box to reflect this uncertainty, guaranteeing that the true body part remains within the adjusted box. Empirical results from a dressing task showed that my

¹Explicit feedback refers to feedback where users directly indicate their preferences, such as choosing between two options.

method generated $\sim 9\text{cm} \times 9\text{cm} \times 9\text{cm}$ boxes that contained the true elbow position 92% of the time. In comparison, baseline methods either produced boxes that contained the elbow only 68% of the time, or up to 81% of the time, but required much larger boxes of $\sim 18\text{cm} \times 15\text{cm} \times 19\text{cm}$. These findings demonstrate that my method, by precisely assessing uncertainty, effectively balances between ensuring that the true body part lies within the box and maintaining a manageable box size for practical use in safe robot motion planning.

#3 Anticipating User Intentions for Efficient Coordination. In multi-step tasks, users' intended sequences of steps are not directly observed by the robot. Estimating these intentions is essential for robots to coordinate safely and efficiently. Prior work often assumed that a user's task step was fixed or followed a linear path. However, in real-world scenarios, a step can branch into multiple next steps. For example, as shown in the figure on the right, after grabbing sandwich ingredients, a user might choose to assemble it at either station 1 or 2 with different sauce options.



My collaborator and I developed the first framework that enables robots to coordinate with users' branching intended steps. This framework estimates user intentions from implicit feedback, such as hand movements, and adapts the robot's actions to complete the task while subtly prompting user responses to refine its understanding of user intentions [4]. We later extended this framework to include verbal communication, enabling robots to use spoken prompts to better understand users' intended steps [5]. Verbal prompts improved task performance by up to 24%, based on a metric balancing task efficiency and communication cost. This work later became a permanent exhibit at the MIT Museum, showcasing its adaptability to human intentions in the real world.

FUTURE RESEARCH

My future research aims to enable robots to deliver personalized assistance that enhances human quality of life, providing both physical support (e.g., dressing, walking) and cognitive support (e.g., reminders, motivation for exercise and social engagement, and tutoring). I envision designing solutions for a wide range of users, including older adults, individuals with cognitive impairments, and children. Methodologically, I seek to bridge human-centered research and algorithmic development to advance personalized robotics. Specifically, I plan to create methods that allow robots to use human implicit feedback to uncover cognitive and emotional processes, effectively learn user preferences, and deliver comfortable, adaptive assistance. Additionally, I aim to develop methods that leverage data from multiple users to enhance personalization for diverse individuals.

#1 Understanding User Cognitive Processes for Advanced Preference Learning. Understanding user preferences is crucial. Building on my research contribution #1, I aim to enhance preference learning by enabling robots to interpret users' cognitive processes from both human explicit feedback, such as choices, ratings, and rankings, and implicit feedback, such as response times and eye movements. For instance, after users receive assistance from a robot with dressing and provide ratings on their overall experience, their ratings may not tell the full story. One user might appreciate the initial part of the assistance but find the latter part unsatisfactory, while another may perceive the entire task as average, yet both could give similar overall ratings. Capturing these underlying cognitive processes can help robots differentiate such cases and learn user preferences more effectively. A key challenge lies in understanding how these cognitive processes evolve in response to the robot's actions over time and ultimately shape the explicit feedback users provide. Traditional models, such as the Drift-Diffusion Model [10] used in my previous work, assume static cognitive processes that lead to simple choices, making them unsuitable for prolonged interactions and richer explicit feedback. To address this, I plan to collaborate with psychologists and cognitive scientists to develop models that capture evolving cognitive processes and design algorithms that use implicit feedback to enhance preference learning.

#2 Decoding User Emotions for Comfortable Assistance. Considering human emotion is essential for making robot assistance adaptive and comfortable. For instance, a robot could monitor a user's sweat rate to assess their sense of safety and adjust its actions accordingly. I plan to collaborate with experts in wearable sensing to

tackle the challenge of interpreting the unobserved, evolving emotional states from human implicit feedback, such as physiological signals. Since emotions cannot be directly measured, researchers rely on explicit feedback, like survey ratings, to train models that interpret emotions from implicit feedback. Current methods typically involve collecting implicit feedback during interactions and obtaining explicit feedback afterward to train these models. However, the predetermined timing between explicit and implicit feedback may not effectively reduce a robot's uncertainty about user emotions. Moreover, current approaches often use offline learning, where robots first gather data, train models, and then apply those models during assistance. This creates a "chicken-and-egg" dilemma: robots need reliable models to assist effectively but require interaction data to build those models. To address these challenges, I will build on my research contributions #1 in online preference learning and #3 in decision-making with unobserved and evolving human intentions. I plan to develop online learning algorithms tailored for unobserved, evolving emotional states, enabling robots to assist while continuously gathering implicit feedback, requesting explicit feedback as needed, and adapting their assistance to enhance user comfort.

#3 Scaling Up Personalization through Collaborative Learning. Personalizing robot assistance and understanding human behavior, cognition, and emotions require substantial user data. However, individual user data is often limited. Inspired by recommender systems and Internet of Things, I aim to scale my past research from learning for one user to leveraging data from multiple users to enhance personalization. A key challenge is that implicit feedback is often sensitive, raising privacy concerns with centralized data storage. Personalized federated learning offers a promising solution, allowing models to be trained locally on user devices while sharing insights across devices without transferring raw data, preserving privacy while enhancing personalization. Adapting federated learning for assistive robotics is particularly challenging because user robots vary in design, tasks, environments, and safety requirements. I aim to address these complexities and develop new solutions that ensure efficient, privacy-preserving personalization in assistive robotics.

References (* denotes equal contribution or alphabetical ordering)

- [1] **Shen Li***, Yuyang Zhang*, Zhaolin Ren, Claire Liang, Na Li, and Julie A Shah. Enhancing preference-based linear bandits via human response time. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2024. **Oral presentation**. [PDF].
- [2] **Shen Li**, Nadia Figueroa, Ankit Shah, and Julie A. Shah. Provably Safe and Efficient Motion Planning with Uncertain Human Dynamics. In *Proceedings of the Robotics: Science and Systems (R:SS)*, 2021. [PDF].
- [3] **Shen Li***, Theodoros Stouraitis*, Michael Gienger, Sethu Vijayakumar, and Julie A. Shah. Set-based state estimation with probabilistic consistency guarantee under epistemic uncertainty. *IEEE Robotics and Automation Letters (RA-L)*, 2022. [PDF].
- [4] Vaibhav V. Unhelkar*, **Shen Li***, and Julie A. Shah. Semi-supervised learning of decision-making models for human-robot collaboration. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2020. **Oral presentation, acceptance rate: 5%**. [PDF].
- [5] Vaibhav V. Unhelkar*, **Shen Li***, and Julie A. Shah. Decision-making for bidirectional communication in sequential human-robot collaborative tasks. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2020. [PDF].
- [6] Tariq Iqbal, **Shen Li**, Christopher Fourie, Bradley Hayes, and Julie A. Shah. Fast online segmentation of activities from partial trajectories. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2019. [PDF].
- [7] Yanwei Wang, Nadia Figueroa, **Shen Li**, Ankit Shah, and Julie Shah. Temporal logic imitation: Learning plan-satisficing motion policies from demonstrations. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2023. **Best Student Paper**. [PDF].
- [8] Ankit Shah, **Shen Li**, and Julie Shah. Planning with uncertain specifications (puns). *IEEE Robotics and Automation Letters (RA-L)*, 2020. [PDF].
- [9] Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 1952. [PDF].
- [10] Roger Ratcliff. A theory of memory retrieval. *Psychological review*, 1978. [PDF].