

Methods for Robot Behavior Adaptation for Cognitive Neurorehabilitation

Alyssa Kubota and Laurel D. Riek

Department of Computer Science and Engineering, University of California, San Diego,
La Jolla, California, USA; email: akubota@ucsd.edu, lriek@ucsd.edu

Annu. Rev. Control Robot. Auton. Syst. 2022.
5:109–35

First published as a Review in Advance on
September 27, 2021

The *Annual Review of Control, Robotics, and
Autonomous Systems* is online at
control.annualreviews.org

<https://doi.org/10.1146/annurev-control-042920-093225>

Copyright © 2022 by Annual Reviews.
All rights reserved

Keywords

human–robot interaction, behavior adaptation, neurorehabilitation,
healthcare robotics, cognitive robotics, social robotics

Abstract

An estimated 11% of adults report experiencing some form of cognitive decline, which may be associated with conditions such as stroke or dementia and can impact their memory, cognition, behavior, and physical abilities. While there are no known pharmacological treatments for many of these conditions, behavioral treatments such as cognitive training can prolong the independence of people with cognitive impairments. These treatments teach metacognitive strategies to compensate for memory difficulties in their everyday lives. Personalizing these treatments to suit the preferences and goals of an individual is critical to improving their engagement and sustainment, as well as maximizing the treatment's effectiveness. Robots have great potential to facilitate these training regimens and support people with cognitive impairments, their caregivers, and clinicians. This article examines how robots can adapt their behavior to be personalized to an individual in the context of cognitive neurorehabilitation. We provide an overview of existing robots being used to support neurorehabilitation and identify key principles for working in this space. We then examine state-of-the-art technical approaches for enabling longitudinal behavioral adaptation. To conclude, we discuss our recent work on enabling social robots to automatically adapt their behavior and explore open challenges for longitudinal behavior adaptation. This work will help guide the robotics community as it continues to provide more engaging, effective, and personalized interactions between people and robots.

ANNUAL
REVIEWS **CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

1. INTRODUCTION

Approximately 15–20% of the world's population has a disability that greatly impacts their independence (1). Such disabilities can affect their ability to perform activities of daily living (ADLs) (e.g., eating and bathing) and instrumental ADLs (e.g., managing medication and finances) (2, 3). However, the number of people who need support exceeds the availability and resources of full-time care providers; informal caregivers (e.g., family) must often assume much of the care responsibility (4, 5), yet they are provided with few resources to do so, leading to stress and burnout (3).

Robots have shown great potential to help people across numerous aspects of health and wellness. Examples range across many settings, including homes, clinics, and hospitals, and different tasks, including reducing clinician workload, supporting people with disabilities, and supporting caregivers (2, 3, 6–12).

Robots can enable clinicians to have more meaningful and productive interactions with people even if there is reduced face-to-face interaction overall, such as during the COVID-19 pandemic. They have the potential to enable clinicians to treat more patients, particularly if the robots are deployed longitudinally in a person's home to help observe, assist with ADLs, or extend interventions. Additionally, robots can lower the cost of treatment for patients, as their use reduces the need for the clinician's time (13). Robots also have potential to provide support to people who live in areas where access to clinicians is limited or nonexistent (e.g., rural areas) and possibly reduce health disparities (14).

Across both the research and commercial sectors, socially assistive robots, which provide assistance through social interactions, are being deployed in people's homes to support their health and prolong their independence. **Figure 1** shows several examples, which vary in form and function. For example, PARO is a zoomorphic, pet-like robot that has been shown to help reduce negative feelings such as stress and anxiety among people with dementia (PwD) and their caregivers (6, 21) and can also alleviate pain and improve mood (22). Researchers are also exploring socially assistive robots to help people with cognitive impairments learn to manage their conditions through cognitive training (11, 19, 23, 24).

In cognitive training and other behavioral treatments, it is critical to personalize training to individuals' preferences, needs, and goals to maximize its applicability to their lives. Personalization helps improve engagement with training, retention of material, and long-term adherence to training (25, 26). These treatments are traditionally led by a human neuropsychologist or



Figure 1

Examples of robots used to support people with cognitive impairments. These robots vary widely in their morphologies, which include mobile, tabletop, humanoid, mechanistic, and zoomorphic. From left to right: Bandit (15) (provided by Maja Matarić), Care-O-bot (16, 17) (provided by Fraunhofer IPA; <https://www.care-o-bot.de/en/>), KOMPAI-2 (18) (provided by Kompai Robotics; <https://kompai.com>), Kuri (19) (provided by Mayfield Robotics), Mabu (20) (provided by Catalia Health), and PARO (6) (provided by Carlton SooHoo, Panospin360; <https://www.panospin360.com>).

cognitive therapist who works closely with a person to determine their needs and goals and tailor training to them.

Recently, researchers have explored how to enable robots to help facilitate training while supporting clinicians and caregivers (3, 19, 23, 24, 27–30). However, how a robot can facilitate training and adapt its behavior in response to a user's behavior and preferences is still an open challenge. These behaviors can consist of physical movements and social interaction strategies (e.g., initiative and personality), and preferred behaviors may vary across cultural and personal backgrounds (31, 32). Managing behaviors becomes particularly challenging when working with a progressive condition such as dementia, as a person's preferences, cognitive abilities, and moods may change quickly during training (19).

Longitudinal behavior adaptation methods from human–robot interaction (HRI) can help address these challenges. These approaches enable a robot to create a model of a person (e.g., their personality and preferences) (31) to guide how it interacts with and responds to them. Doing this accurately and consistently is critical when working with people with neurodegenerative conditions, such as PwD (19, 33).

In this review, we explore longitudinal adaptation methods that can enable social robots to personalize their behavior to individuals, thereby improving a person's longitudinal engagement and adherence to training. We explore dementia as a specific exemplar context for cognitive neurorehabilitation. First, we provide an overview of the application domain of neurorehabilitation (Section 2) and review several existing robots used to support PwD (Section 3). Next, we outline key principles for researchers to be mindful of as they design robots for neurorehabilitation (Section 4). We then examine some common communication modalities (Section 5) and explore technical approaches to behavior adaptation (Section 6). To conclude, we explore some open problems in this area (Section 7) and discuss our recent work on robot behavior adaptation (Section 8).

As people begin to integrate robots into their lives, robots must sustain engagement over long periods of time to maximize each interaction's efficacy. This work will guide the robotics community to enable robots to personalize their behavior to an individual's needs and preferences, so that they can more effectively help people accomplish a multitude of long-term goals, from overcoming memory challenges to living healthier lifestyles.

2. MILD COGNITIVE IMPAIRMENT AND NEUROREHABILITATION

2.1. Mild Cognitive Impairment

Dementia is an irreversible syndrome that entails a noticeable decline of cognitive function (3, 34). Approximately 11% of people over age 65 are impacted by dementia, and each case is unique. Symptoms can range across the spectrum, from forgetfulness in the early stage to difficulty recognizing friends and family in the late stage. Dementia can affect a person's physical abilities, mental abilities, and behavior and can lead to hazardous behaviors such as wandering, medication errors, and domestic or financial abuse. There are no known cures to slow or prevent its onset, which can reduce quality of life for family members when they adopt the role of informal caregivers (35).

Mild cognitive impairment (MCI) is the prodromal, or intermediate, state between normal aging and several neurodegenerative disorders, such as Alzheimer's disease and vascular dementia (26, 36). An estimated 20% of adults over age 65 experience MCI, approximately 10% of whom convert to some type of dementia each year (26, 37). To date, no existing pharmacological treatments have proven effective for slowing or preventing this conversion, but studies suggest that behavioral treatments can help (26).

MCI can affect numerous cognitive domains, including memory, visuospatial functioning, complex attention, and executive functions, though not to a level of severity that would warrant a

diagnosis of dementia (36, 38). Studies indicate that many people will remain at the MCI stage without ever converting to dementia, and up to 40% of those with MCI will return to normal levels of cognitive functioning over time (26). However, as people lose their independence, MCI can severely impact their quality of life (26, 39). It can also adversely affect their family members, put strain on their relationships with the person with MCI, and cause stress (3, 26, 40–42). This change in lifestyle and role can cause feelings of guilt, anxiety, and depression in a person with MCI and their caregivers (40, 41).

2.2. Neurorehabilitation

Many researchers have explored strategies to promote the reablement of PwD or mitigate the impact of dementia on their function to promote independence (43). In particular, nonpharmacological approaches such as behavioral therapy can slow the onset of MCI, which can prolong independence and maintain quality of life (26). Treatment approaches include cognitive rehabilitation and restoration therapies, which aim to minimize or compensate for lost cognitive function in everyday life. Among the most widely used strategies are compensatory cognitive training and restorative cognitive training (26).

Compensatory cognitive training teaches a person with MCI metacognitive strategies to help bypass impaired function and minimize its impact on daily life (26, 39). These strategies may include reorganizing their environment (e.g., always placing their keys next to the door when they return home), integrating new tools into their daily routine (e.g., routinely keeping and checking a daily planner), and using different skills to compensate for memory difficulties (e.g., using visual imagery or acronyms). Depending on an individual's impacted cognitive abilities, clinicians may prescribe different training regimens to focus on specific skills. Compensatory cognitive training has been shown to improve cognitive performance and daily functioning in people with MCI, and these improvements are often sustained even after a person has completed training (26). In our work, we focus on employing compensatory cognitive training with a robot (19).

Restorative cognitive training, by contrast, attempts to enhance or restore a person's lost cognitive abilities. It relies on consistent practice and repetition of standardized cognitive exercises designed to target specific skills, such as attention or memory (e.g., drill and practice). While this approach can help strengthen neural circuits and improve a person's performance on similar tasks, these exercises are generally standardized (i.e., not personalized to an individual) and may not be relevant to a person's everyday life (44). Furthermore, these skills typically do not generalize well (i.e., transfer) to other tasks (26).

3. ROBOTS FOR NEUROREHABILITATION

Robots for physical neurorehabilitation typically help people by physically supporting or correcting movement with the goal of restoring neuromotor function—for example, restoring or supplementing limb function in people who had a spinal cord injury or stroke (7–9). These robots take many forms, such as robotic arms to help people control their arms and hands to complete ADL tasks (7, 8) and exoskeletons to help people walk (9).

Researchers also use socially assistive robots to support cognitive neurorehabilitation. (To our knowledge, there are presently no commercially available robot systems to administer cognitive neurorehabilitation.) These robots interact with people through social signals such as speech or gestures. They help people practice cognitive skills and social interactions that they can transfer to everyday life (15, 45).

In addition to dementia, researchers have increasingly explored the use of socially assistive robots to support people with social and developmental disorders, particularly children with

autism spectrum disorder or attention-deficit/hyperactivity disorder (46–49) and people with schizophrenia (12, 50, 51). For instance, children with autism spectrum disorder expressed more spontaneous behavior, both nonverbal and emotional, after interacting with a robot mediator, which they were able to translate to interactions with another person (52, 53). Robots can also help improve communication between older adults with schizophrenia and their medical providers and increase their engagement with recreational activities (12, 50).

3.1. Benefits of Robots for Neurorehabilitation

Robots present many exciting opportunities for supporting rehabilitation. They are a natural fit for the repetitive, task-oriented nature of many cognitive interventions, such as restorative cognitive training exercises, which are often structured. They can also provide real-time, adaptive feedback, providing unique opportunities for rehabilitative therapy.

While computer-assisted strategies for administering neurorehabilitation exercises have been shown to improve attention, memory, and executive skills in people with memory impairments (25), robots have even greater potential to improve training, as their physical embodiment plays an important role in stroke patient compliance and engagement (33, 54). Robots can increase engagement and enjoyment in social interactions due to their increased capacity for rich communication as compared with virtual systems (55). They have many attributes that are important for initiating and sustaining interactions, including shared physical context, physical movement, and the ability to appear to be observing a user (56).

A robot can also monitor and assess a person's well-being or task behaviors, which can be shared with their care team as well as with the user. For example, during cognitive training, a robot could collect information on task performance and progress. It may also infer other attributes, such as their level of engagement and interest through gaze tracking, proxemics, or voice recognition. The information that a robot gathers has the potential to provide clinical insights that may help reduce a clinician's cognitive load. Clinicians can use this information to adjust training to match a person's abilities and preferences. They may also use it to help inform a person about their condition or to understand what aspects of the training are most effective. Section 5 provides an overview of various behaviors about a person that a robot can sense, as well as how the robot may respond to those behaviors.

3.2. Exemplar Robots for Mild Cognitive Impairment and Dementia

There are many robots to support PwD (see **Figure 1**). They fill numerous roles, such as assistive robots to help users complete ADLs, companion robots for emotional support, or robots to facilitate therapy or coach people practicing cognitive skills. **Figure 2** provides an overview of selected robots, some of which were designed specifically for PwD and others of which have been applied in this space.

These robots are typically mobile and provide monitoring and care, helping to ease the responsibilities of informal caregivers (e.g., family and friends) and extending the independence of PwD. Their capabilities may include reminding a person to take medication, facilitating communication between the person and their care network (e.g., video calls with clinicians or family), and delivering cognitive stimulation (57, 58). Others include walking assistance, fetching items, or setting a table for people with mobility difficulties (16). They are usually used within a home, though some researchers are also exploring robots that can accompany users on errands outside of the home (10, 63).

Robots may also serve as companions for people with MCI and dementia. Many of these robots have been shown to reduce stress and anxiety while improving relaxation and motivation among

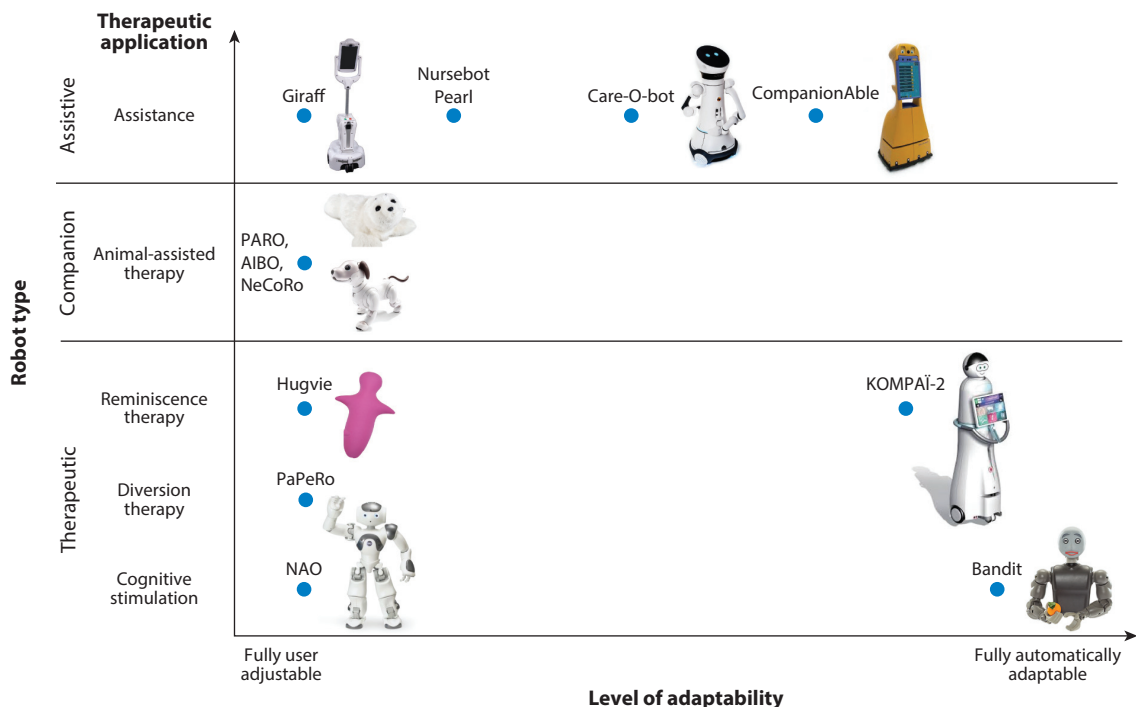


Figure 2 Exemplar robots that have been used to support neurorehabilitation and therapy. Specifically for dementia, robots are typically used for assistance, companionship, or therapeutic applications such as animal-assisted therapy or reminiscence therapy. Many are fully user adjustable, while others can automatically adapt their behavior in response to users. Their morphologies can vary depending on the application, such as mobile robots used to provide physical assistance or tabletop robots used for cognitive therapy. From left to right and top to bottom, the images are as follows: Giraff (57) (provided by Camanio AB), Care-O-bot (16, 17) (provided by Fraunhofer IPA; <https://www.care-o-bot.de/en/>), CompanionAble (58) (provided by Steffen Müller), PARO (6) (provided by Carlton SooHoo, Panospin360; <https://www.panospin360.com>), AIBO (59) (provided by Sony), Hugvie (60) (provided by Hiroshi Ishiguro Laboratories, Advanced Telecommunications Research Institute International), KOMPAI-2 (18) (provided by Kompai Robotics; <https://kompai.com>), NAO (11) (provided by SoftBank Robotics), and Bandit (15) (provided by Maja Matarić). Images are not shown for Nursebot Pearl (10), NeCoRo (61), and PaPeRo (62).

PwD and their caregivers (6, 21). They can help stimulate interaction and serve as a point of connection between PwD and their caregivers (6). Many of these robots resemble animals, making them recognizable even to people with severe memory impairments. For instance, PARO (6) is based on a baby harp seal, and AIBO (59) resembles a dog. These types of robots do not necessarily communicate with people via speech, but can instead move or make sounds in response to stimuli such as touch, sound, or light (6).

These companion robots are often used in therapy. For instance, many of the aforementioned robots serve as safer alternatives to real animals in animal-assisted therapy and activities, often in hospitals and nursing homes (6, 59, 61). In addition, researchers have explored using PARO to facilitate multisensory behavior therapy (64, 65), which stimulates different senses in a controlled setting to reduce agitation in uncontrolled ones.

More recently, researchers have used robots to facilitate reminiscence therapy among PwD (18, 60). Reminiscence therapy aims to help people recall long-term autobiographical memories with the aid of photographs, music, familiar objects, and so on. It is highly regarded by participants and therapists and viewed as enjoyable and effective (66). The approach is generally conversational,

guided by either a human therapist or the robot itself, using a microphone and speaker in the robot to communicate with the person (18, 60). In robot-guided sessions, a robot relies on user-specific knowledge (e.g., photos from an event or a favorite location) to prompt the user and maintain conversation and memory recollection.

Another role that robots may take for MCI and dementia is that of a coach. These robots are often used to facilitate and assist with restorative cognitive training exercises. For example, the Bandit robot plays cognitive stimulation games with users and adjusts the difficulty based on their performance (23). Similarly, researchers have programmed humanoid robots such as NAO for clinicians to use to assist with memory training programs (11).

4. PRINCIPLES FOR DESIGNING NEUROREHABILITATION TECHNOLOGY

When designing technology for people with disabilities, one should consider three key factors to ensure that it is usable and acceptable: personalization, adaptation, and inclusion (12, 67). Personalization refers to tailoring the system to an individual by considering factors such as their needs, goals, or preferences. Adaptation is the ability of a technological system to automatically modify its behavior to be personalized to an individual. Inclusion means involving stakeholders throughout the process of developing technology, particularly the intended users of that technology. These considerations are particularly important to prevent unexpected consequences for potentially vulnerable populations, such as the exacerbation of disability-based bias (68, 69).

4.1. Personalization and Adaptation

It is critical that neurorehabilitative technologies be personalized to an individual, from simply including their name to adapting to suit their unique preferences and abilities (70). This is important when developing any technology for people who may not be represented by a “typical” user (67). In fact, early studies on the efficacy of cognition-based interventions suggested that they were ineffective and inappropriate for people at risk of cognitive impairments because they could provoke frustration and depression in both a person and their caregivers (40, 71), which was likely a result of their reliance on repetition and structure (e.g., memorizing and repeating a specific list of words) without any clear connection to the individual’s life, abilities, or interests.

Especially when developing technology interventions for health contexts, each individual has unique circumstances that can significantly impact how they interact with the technology. For instance, up to 77% of older adults with MCI may be managing comorbidities (e.g., MCI and diabetes) or have different living situations (e.g., living alone or in a nursing home) (72). A system that is not personalized may cause needless stress or frustration for the user and their caregivers or have other detrimental effects on their health. By tailoring the training to an individual and meeting them where they are in terms of their performance and abilities, modern neurocognitive interventions become significantly more effective and beneficial for people with cognitive impairments and their caregivers as compared with nonpersonalized interventions (27, 39).

The ability for technology to be personalized calls for the system to be adjustable by a human, adapt its own behavior, or both. There are many situations in which a clinician, caregiver, or user may want to control or adjust a robot’s behavior. For instance, with the domain expertise from clinicians and the fundamental personal knowledge from caregivers and users, they may already have a good idea of how they want a robot to behave to facilitate and complement the training. Additionally, clinicians and other users may want to modify the system to reflect the training. In a home setting, a user or caregiver may want to adjust behavior without the help of a clinician. Thus, any mechanism to manually adjust the system should be easily learnable and usable by all stakeholders.

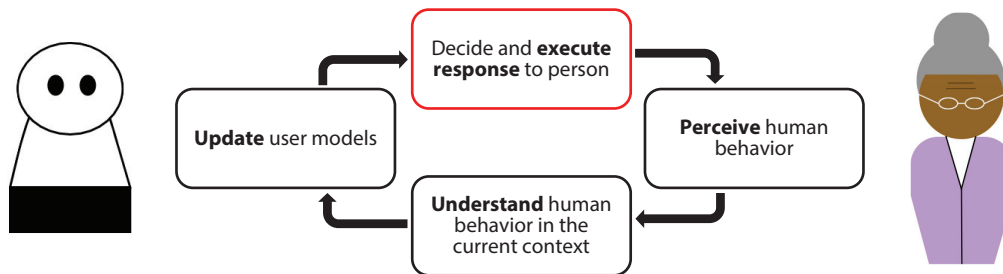


Figure 3

The general framework of systems that adapt to users. In this review, we focus on deciding and executing the robot's response to a person. Figure inspired by Reference 74.

There are also situations in which it may be beneficial to automatically adapt to a user. Conditions such as dementia can be progressive, and the person receiving training may be undergoing cognitive changes at a pace that is difficult for others to keep up with. A robot that uses a computational model for automatic adaptation may have the advantage of learning and remembering information about a user more quickly and accurately than a person can.

Automatic adaptation alleviates the responsibility of caregivers or clinicians to continually adjust a robot's behavior, enabling them to then spend more time in face-to-face interactions with an individual. Additionally, studies indicate that older adults prefer assistive systems that allow them to control the system while still being adaptable over fully adjustable ones (73). Thus, automatic adaptation to a user can lead to more rapid adjustments to a training regimen, which may improve its efficacy and sustain a user's engagement.

To automatically adapt to a user, a system must be able to perceive and interpret a person's actions and respond in a meaningful way (see **Figure 3**). This involves considering what a robot will sense about a user and how to obtain those data. For instance, what sensors will it need and where will they be placed? What information will it infer implicitly (e.g., from sensor data or observations) as opposed to obtaining explicitly (e.g., through questionnaires or surveys)? Section 5 provides an overview of potential sensing modalities and inputs.

Once a robot has this information, it needs to contextualize and understand what it means about a person. This could be their current state (e.g., mood or task performance) or an overarching understanding of the person (e.g., ability level). Finally, robots need to know how to modify their behavior and respond to a user. Roboticists employ numerous computational models to achieve this, which we discuss further in Section 6.

4.2. Inclusive Design

It is important for roboticists and researchers interested in building assistive robots to involve stakeholders throughout the development process. This is exceptionally true while developing a robot to be deployed in a person's home with the goal of supporting their health. These stakeholders may include the primary robot user, their healthcare providers, and their caregivers, who may or may not be living with them (2, 3, 42).

Nilhil de nobis, sine nobis, or "nothing about us without us," is a prominent motto of disability activists (75). It conveys that people with disabilities themselves know what is best for them and that they are integral in any conversation that may affect their life and community. In other words, they must be consulted regularly throughout the technology development process, from conceptualization to testing. As roboticists often develop technology for conditions they have no

personal experience with, involving people with disabilities early and often will help avoid making assumptions about the community's goals, ensure their needs are met, and help empower them. This will ensure the maximum utility, usability, and acceptance of the technology by users as well as other stakeholders (2, 76).

When codesigning technology with stakeholders, it is important to be transparent about what the technology is capable of. As there are no known approaches to significantly impact the course of dementia, technology should encourage stakeholders to “live well with dementia” (43, p. 455). This means setting realistic expectations about the benefits stakeholders can expect from the technology—for instance, how it could change the roles of clinicians and caregivers, the extent of its impact on a user's training process and results, or the data it collects. The onset of the condition being treated is possibly one of the most challenging experiences the stakeholders have undergone, so developers must develop their trust and maintain compassion for them throughout the development process.

Additionally, those receiving neurorehabilitation are likely vulnerable populations and do not necessarily have the technological literacy to effectively operate a system. Low technological literacy and cognitive impairment can also affect informed consent (3, 12, 77). Developers of this technology must be mindful of this and work closely with experts in these communities to protect user privacy while maintaining the system's utility.

5. SENSING AND RESPONDING TO HUMAN BEHAVIOR

Modifying robot behavior to be personalized to an individual is crucial for maintaining engagement and ensuring the efficacy of the system, particularly for health applications (78, 79). In order for a robot to effectively adapt its behavior to a user, it must perceive the user's actions and behavior, understand what those mean in the given context, and respond accordingly (33, 80). Below, we identify some features about people that robots can sense as well as behaviors that robots can modify in order to personalize interactions.

5.1. Perceiving and Understanding Human Behavior

Throughout an interaction, there are many ways robots can learn user preferences and abilities. One approach is to first perceive a person's low-level behavior, then infer how those behaviors translate to higher-level attributes. Robots can gather this low-level information via the use of sensors or through interaction or performance data collected by the system. Examples of low-level behaviors that a robot may gather include the user's speech (e.g., what they say and how they say it), gestures and movement (e.g., human activity recognition), and physiological signals (e.g., heart or respiration rate). Performance data are typically application specific and depend on the task(s) (e.g., accuracy or time to complete a task).

Some major factors to consider when choosing which sensors to use are what kinds of sensors a robot already has, whether others can be easily placed in the environment, and what kind of information would be worthwhile to collect, process, and possibly store. These sensors may be on a robot, placed in the environment, or worn by a person. For instance, cameras or microphones may be mounted on a robot or in the environment depending on the context, while physiological or inertial sensors are typically worn by a person.

These sensor and interaction data are relatively low level and can be used to infer higher-level information about a user's state or preferences (31). For instance, robots can use data they acquire from RGB-D cameras to track a person's gaze or movements, then use these features to infer higher-level features, such as how engaged or bored they are (e.g., the person is likely to be engaged if they maintain eye contact with the robot and gesture often).

Table 1 Common features that robots can sense about people using a variety of sensors or indicators; low-level interaction data can be used to infer higher-level information about a user's state

Behavior perception	Common sensor(s)/indicator(s)	Description
Low level		
Speech/prosody	Microphone	Speech is a common means of communicating with a robot. In addition to what a person is saying, their prosody and tone may convey important information.
Gesture/movement	RGB camera, motion capture, gyroscope, accelerometer	Arm and hand gestures are a common means of communicating with a robot, both implicitly (e.g., everyday activities) and explicitly (e.g., specific gestural commands).
Eye contact/gaze	Infrared camera, RGB camera	Gaze tracking helps determine where a person is looking.
Touch	Capacitive touch sensor, force sensor, pressure sensor, strain gauge, switches	Determining whether a person is touching a robot or where they are touching can add realism to interactions.
Physiological signals	EEG sensor, EMG sensor, heart rate monitor, respiration sensor, thermometer	Signals generated by a person's body, usually acquired from specialized wearable sensors, can help determine their state.
Explicit feedback	Questionnaires, surveys	Asking users for their input directly is a straightforward way to obtain information.
Task performance	Application specific	In neurorehabilitation, the robot may administer cognitive training activities with quantifiable scores.
High level		
Engagement	Eye contact, touch, speech	Longer and more positive interactions with a robot can help sustain interactions over longer periods of time.
Mood	Physiological signals, speech	A user's current mood can help inform how a robot should best interact with them.
Motor abilities	Touch, movement	A user's motor abilities can help inform their preferred means of communicating with a robot. For instance, a user with tremors may prefer speaking instead of pressing buttons.
Cognitive abilities	Task performance, speech	A user's cognitive abilities can influence their goals and what treatment regimens may be most effective.

Abbreviations: EEG, electroencephalography; EMG, electromyography.

An alternate approach is to ask a user about their preferences—for example, in a questionnaire or survey (31). This is a straightforward and direct means of obtaining information that does not require additional sensors. However, it risks people providing their ideal answers rather than completely truthful ones. **Table 1** provides an overview of common features that social robots for neurorehabilitation can sense about people, at both low and high levels.

Once a robot perceives a person's behavior, the robot must consider how that behavior relates to the person's current state and/or their overarching condition. How a robot interprets a person's behavior may depend on the application, length of interaction, or other circumstances. It is important for the robot to understand a person's actions and their current state (e.g., mood) in order to maintain natural, real-time interactions. For example, a robot may use a person's body language or task performance to infer whether the person is frustrated with or challenged by a cognitive training exercise (33).

Particularly over long-term interactions, such as while completing a cognitive neurorehabilitation session, it is important for the robot to store and update a model of a person, including their preferences, needs, and abilities (31). A robot can use this model to understand what

behavior is typical for a person, track their progress over time, and recognize if they deviate from what is expected (e.g., recognizing if the person is more agitated or more forgetful than usual). Both understanding individual actions and translating them into a more thorough model of a person are important for personalizing robot behavior.

5.2. Synthesizing Robot Behavior in Response to People

Effective HRI requires that robots understand people and respond to them. Individual robot actions can be guided by a fundamental model of its interaction style (e.g., personality and role) (78, 79). In the context of neurorehabilitation, these behaviors can help improve a user's enjoyment of a training regimen and thus its efficacy (19, 29, 31).

When interacting with people, a robot can personalize its behavior in response to a person in numerous ways. At a low level, movement, speech, and visual cues are some major ways a robot can communicate. Movement generally consists of physical motion of the base or limbs. Speech can include dialogue, speed, prosody, tone, or other sounds. Visual cues may be a change of expression, text or images on a tablet, or other cues.

A robot may change its communication modalities based on user abilities or state. For example, a user with tremors may prefer to communicate via speech, whereas someone who is nonverbal may prefer a tablet interface. Depending on a robot's capabilities, it may also change its effectors or display to convey emotion or emphasis to enhance an interaction.

These low-level behaviors can be utilized to produce higher-level aspects of the interaction that consider user preferences and needs. By immediately reacting to a person, a robot can create more natural and engaging interactions, such as by maintaining eye contact during conversation. For instance, if a user seems distracted, a robot may change its dialogue and tone to return their attention to the robot. This can help maintain engagement throughout an interaction, thus improving retention of material and overall enjoyment (81).

Similarly, a robot may change longer-term aspects, such as its personality. For example, if a person responds better to an encouraging personality than an assertive one, the robot can provide more encouragement throughout training. In this way, a robot can update its model of a person and use it to guide the interactions, modifying its behavior to be more personalized to an individual. This can help maximize a person's adherence to a training regimen and improve their perceptions of the robot (29, 31). **Table 2** provides an overview of some common social robot behaviors that may be altered throughout interactions with a person.

6. COMMON TECHNICAL APPROACHES TO BEHAVIOR ADAPTATION

A key element for enabling robots to adapt their behavior to a user is understanding how the data they receive can inform their actions, as well as how a user responds to those actions. There are countless computational methods researchers have used to imbue social robots with this ability, both within and outside of the context of neurorehabilitation. **Table 3** provides a summary of common approaches, which are further discussed below. (As many social behaviors are not robot specific—e.g., dialogue—we also include select systems that were demonstrated on non-physically embodied systems, such as virtual agents.)

While perceiving and understanding human behavior are important aspects of knowing how a robot should respond, the area of human behavior analysis is vast, and approaches may vary widely depending on the behavior being perceived. As this article focuses on methods for robot behavior adaptation, we discuss approaches that assume the human behavior is already recognized, as well as those that embed human perception into their process. For detailed surveys of human behavior analysis, see References 112 and 113.

Table 2 Common social robot behaviors that may be altered throughout interactions with a person; robots can modify low-level behaviors to personalize high-level aspects of an interaction in order to accommodate a user’s preferences and needs

Modality	Description
Low level	
Movement/speed	Movement can be used for mobility, communication, and to help interactions feel natural.
Speech/speed/prosody/sounds	Speech is a common means of communication for both humans and robots. In addition to dialogue, a robot may adjust its speed, prosody, or other sounds to improve clarity/function.
Screen display	Robots may use a tablet when communicating with a user.
Facial expressions	Many social robots have faces with dynamic expressions. As people have a tendency to anthropomorphize robots, even those that are not humanoid (82), robots can change their facial expressions to create a more natural and interesting interaction.
Proxemics	Proxemics is the division of physical space around an agent (classified as intimate, personal, social, or public). A robot can control how physically close it is to a person to convey respect or intimacy.
High level	
Personality	A robot may change its personality to suit a user’s preferences. This can be influenced by personal and cultural background; for example, a robot may adopt a more passive communication style in countries where people tend to have more reserved communication styles (83).
Initiative	Initiative is whether a robot initiates interaction with a person or vice versa. This may change with a robot’s role, such as initiating interaction with a user with more severe MCI.
Encouragement	Providing encouragement can help a person be more motivated or less frustrated if they experience trouble with the training regimens.
Personal customization	Integrating personal information into training and therapies can help them be more applicable to a user, improving engagement and efficacy.
Cognitive customization	Adjusting aspects of a training regimen to suit a person’s cognitive abilities can help them practice relevant skills and reduce frustration.
Primary communication modality	Depending on a user’s physical abilities and preferences, a robot may adjust how it receives input from them, such as aural, touch, or visual cues.

Abbreviation: MCI, mild cognitive impairment.

6.1. Finite State Machines

Finite state machines are a relatively straightforward approach to address the behavior adaptation problem (56, 84–86). In a finite state machine, an interaction is broken into states that guide robot behavior. The robot transitions to the next state depending on human and environmental factors.

For instance, Kidd & Breazeal (56) used a finite state machine on Autom, a robotic weight loss coach. A user would engage in a short conversation with Autom once or twice a day. Its dialogue could vary depending on the time of day, time since the last interaction, and recent data input by a user. Each factor filled in parts of the conversation (e.g., Autom said “Good morning” or “Good evening” depending on the time of day). Notably, the robot’s statements also varied depending on the estimated relationship state between the robot and user, which offered a variety of dialogue to avoid repetition during the six-week study.

This approach is useful for short and relatively structured interactions. Many programmers are already familiar with finite state machines, and there are a number of existing libraries to implement them on robots, such as the Python-based SMACH (State Machine) library (114). However, not all interactions can be broken into discrete states, and states are generally defined manually, so implementing long and involved interactions may be infeasible.

Table 3 Common technical approaches for machines that adapt behavior to people

Approach/existing work	Strengths	Limitations
Finite state machines (56, 84–86)	<ul style="list-style-type: none"> ■ Straightforward ■ Existing libraries for implementation on robots (e.g., SMACH) ■ Good for short, structured interactions 	<ul style="list-style-type: none"> ■ Interactions generally cannot be split into discrete states ■ Intricate interactions may be infeasible to implement ■ Does not easily allow for complex behaviors or long-term understanding of a user ■ Does not easily allow for dynamic behavior adaptation
Thresholding (15, 81, 87)	<ul style="list-style-type: none"> ■ Good for reacting to continuous streams of data rather than windows of time 	<ul style="list-style-type: none"> ■ Does not easily allow for complex behaviors or long-term understanding of a user ■ Does not easily allow for dynamic behavior adaptation
Q-Learning RL and variants (MDP) (11, 29, 88–99)	<ul style="list-style-type: none"> ■ Model free, or can learn a model about user behavior/preferences 	<ul style="list-style-type: none"> ■ Assumes the world is fully observable, but a person's preferences cannot always be directly observed ■ Time and storage intensive, which can inhibit real-time interaction ■ Interactions generally cannot be split into discrete states
RL: POMDP (74, 100–103)	<ul style="list-style-type: none"> ■ Model free, or can learn a model about user behavior/preferences ■ Does not assume the world is fully observable, which is beneficial as most human preferences cannot be directly observed 	<ul style="list-style-type: none"> ■ State space becomes intractable for complex interactions ■ Interactions generally cannot be split into discrete states
Hierarchical RL (10, 28, 104, 105)	<ul style="list-style-type: none"> ■ Model free, or can learn a model about user behavior/preferences ■ Makes complex (PO)MDPs more manageable ■ Can handle greater modularity of sensors and behaviors 	<ul style="list-style-type: none"> ■ Does not take combinations of behaviors into consideration and so is not guaranteed to find a globally optimum policy ■ Interactions generally cannot be split into discrete states
Policy gradient RL (33, 106)	<ul style="list-style-type: none"> ■ Naturally handles continuous states and actions 	<ul style="list-style-type: none"> ■ Difficult to derive an appropriate reward function (i.e., preferences from behavior) ■ Need to find appropriate parameter values
Inverse RL (30, 107–109)	<ul style="list-style-type: none"> ■ Learns a reward function from a human expert rather than relying on exploration of different behaviors 	<ul style="list-style-type: none"> ■ Requires feedback from a human expert ■ Human experts do not necessarily behave optimally or rationally
Neural network: MLP (110)	<ul style="list-style-type: none"> ■ Hierarchical layers enable high-level feature extraction from raw or low-level input data 	<ul style="list-style-type: none"> ■ Does not take previous input into account ■ Requires large amounts of training data, which can make it difficult to learn from an individual person ■ Difficult to optimize hyperparameters
Neural network: LSTM (111)	<ul style="list-style-type: none"> ■ Hierarchical layers enable high-level feature extraction from raw or low-level input data ■ Learns temporal features using previous input ■ Reacts to continuous streams of data rather than windows of time 	<ul style="list-style-type: none"> ■ Requires large amounts of training data, which can make it difficult to learn from an individual person ■ Difficult to optimize hyperparameters

Abbreviations: LSTM, long short-term memory; MDP, Markov decision process; MLP, multilayer perceptron; POMDP, partially observable Markov decision process; RL, reinforcement learning; SMACH, State Machine.

6.2. Thresholding

Another approach that roboticists use is thresholding (15, 81, 87). In this approach, a robot receives sensor data from a user and performs an action if the value crosses a given threshold. Tapus et al. (15) used thresholding on a social robot for PwD. It administered a cognitive game and could

adjust the difficulty to improve a person's performance. The robot used an accepted variation band to automatically adjust the difficulty based on the person's performance, with the goal of minimizing reaction time, maximizing the number of correct answers, and maximizing the difficulty level. If the person's performance (i.e., reaction time and correct answers) improved, the difficulty increased, whereas it decreased if they performed poorly. The authors reported increased engagement and improved performance at higher difficulties for PwD, highlighting the importance of adjusting to the user's abilities.

Thresholding is advantageous when dealing with a continuous stream of data and reacting in real time. However, it is best suited for behaviors tied to a specific signal (e.g., increasing voice volume if engagement is low or decreasing task difficulty if performance is low) and does not easily allow for complex reactive behaviors. Additionally, while thresholding enables a robot to react in real time, it would require another underlying control system, such as those discussed below, to support longer-term understanding about a person.

6.3. Reinforcement Learning

In reinforcement learning (RL), an agent learns how to best interact with its environment to maximize its rewards (115). RL configurations are generally represented using a Markov decision process (MDP) defined as $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma)$, where \mathcal{S} is the set of possible states, \mathcal{A} is the set of possible actions the agent can take, \mathcal{T} is the transition probability function between states, \mathcal{R} is the reward function of the environment, and γ is the discount factor for future rewards (115). Actions in MDPs can be deterministic (i.e., performing a given action in a given state always leads to the same next state) or stochastic (i.e., the next state is determined by a probability distribution). The agent aims to learn an optimal policy π , or a mapping of states to actions, that maximizes its expected rewards.

Q-learning is a widely used approach for solving MDPs with unknown reward and transition probability functions. Traditionally, a robot can take an action and observe the associated reward as the environment updates to a new state. Many researchers have applied it to the behavior adaptation problem (11, 29, 88–99). For instance, Tsiakas et al. (29) used it to modify the kind of feedback a robot provided based on a person's engagement in a cognitive training session.

Multiple works frame the behavior adaptation problem as a multiarmed bandit problem (97, 98, 116). The multiarmed bandit problem aims to distribute resources among multiple possible actions with uncertain results in order to maximize the reward, but the current state remains the same. This approach is useful for ensuring that a robot can try each action and observe a person's behavior before relying too heavily on its learned knowledge of the person's reactions to its actions. In behavior adaptation, this can be thought of as selecting behaviors in order to maximize a person's engagement, performance, and so on.

Numerous algorithms exist to help balance exploration of new or uncertain actions with exploitation of existing or learned knowledge, particularly when the available actions have unpredictable outcomes. For instance, Gao et al. (97) implemented Exp3 (Exponential-Weight Algorithm for Exploration and Exploitation) (117) on a Pepper robot for puzzle solving. The robot would learn the person's preference for supportive behaviors (e.g., providing hints and encouragement) and respond to a person's performance (measured by the time since they last made an action, total time elapsed, and correct actions).

Q-learning is model free, meaning it does not require a preexisting model. This is useful for behavior adaptation where human reactions (i.e., rewards) are difficult to define as a model (29, 88). However, Q-learning and MDPs have many limitations, such as assuming the world is fully observable and being time and storage intensive (33, 106, 118). Real-time behavior adaptation for HRI is not always feasible with this approach, as a person's state cannot always be directly observed,

and heavy computations may slow a robot's responses. Thus, there are numerous alternatives to address these problems.

When the world is not fully observable, a partially observable Markov decision process (POMDP) is often more suitable. It is defined as $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \omega, \mathcal{O}, \gamma)$, where \mathcal{S} , \mathcal{A} , \mathcal{T} , \mathcal{R} , and γ are the same as in an MDP; ω is a set of observations; and \mathcal{O} is a set of conditional observation properties (119). The agent does not know its underlying state and must maintain a probability distribution of possible states based on previous observations.

Researchers have used POMDPs for behavior adaptation in health applications such as managing food consumption (74), navigating a robotic wheelchair (100), and helping PwD wash their hands by giving visual or verbal prompts (103). This approach is applicable to behavior adaptation, as a person's state is typically unknown and cannot be explicitly observed by a robot (e.g., a frown could express frustration with training or sadness due to external circumstances). However, the state space can become intractable to manage for complex interactions and multiple behaviors, which may make real-time responses infeasible.

For complex interactions with numerous human and robot behaviors, researchers have used hierarchical RL (10, 28, 104, 105). This approach divides the overall MDP into smaller, more manageable ones, which simplifies the problem and can help reduce memory requirements (28). It can also allow for greater modularity of the system's behaviors; for instance, Chan & Nejat (28) used the MAXQ hierarchical RL approach (120) to abstract their system into a temporal module, state module, and subtask module, each of which considered and controlled specific behaviors in the context of cognitive training. While hierarchical RL approaches can find the optimal policy for each individual MDP, the global policy is not guaranteed to be optimal, as there is no way to consider how behaviors can be combined.

Researchers have also applied policy gradient RL methods for behavior adaptation (33, 106). Policy gradient RL directly adjusts the policy in relation to the gradient to find a locally optimum policy, defined by behavioral parameters that a robot can adjust. It begins with an initial policy that it evaluates according to the reward function. Then, it perturbs the policy by modifying each parameter. Finally, it evaluates the new policy and repeats the process until a local optimum is found.

This approach enables robot learning for continuous states and actions and can update a robot's behavior in real time, both of which are important aspects of behavior adaptation in HRI (106). However, researchers have reported challenges in deriving an appropriate reward function to accurately translate user behavior to explicit preferences (33, 106).

A slightly different approach is inverse RL, in which the robot learns how to behave from an expert agent who is assumed to behave optimally (107, 108). The robot can then use standard RL algorithms following the learned policy to maximize its own reward. In neurorehabilitation, this may entail the robot observing a human therapist guiding the training in order to learn how to respond to a patient in future interactions. However, humans do not always behave optimally or rationally, and it is not always possible to discern an exact policy, so researchers have expanded inverse RL to help overcome these limitations (109).

Additionally, researchers have worked to infer user preferences solely from observing a user, as in observational repeated inverse reinforcement learning (ORIRL) (30). In ORIRL, a robot learns a user's preferences by watching them complete different tasks, then leverages those learned preferences when inferring preferences for future activities.

6.4. Artificial Neural Networks

Recently, researchers have begun to leverage advances in neural networks and deep learning for robot behavior adaptation (110, 111). Neural networks are a broad set of algorithms inspired by

biological neural networks that enable agents to recognize patterns in data, generally without having to define underlying task-specific rules. They have a hierarchical structure where the neurons (a computational unit) of each layer can extract information from the previous one to learn higher-level features. Thus, deep learning approaches with multiple hidden layers have gained popularity for their ability to extract features from raw data without the need for human-defined features, a large source of variation in other learning methods (121).

Senft et al. (110) used a multilayer perceptron (MLP) to enable a robot to learn from a therapist how to interact with children with autism spectrum disorder. An MLP is a supervised feed-forward neural network composed of multiple perceptrons, or binary classifiers, with a unique set of weights (122). The use of multiple perceptrons allows the MLP to approximate nonlinear functions for multiclass classification. The MLP used by Senft et al. (110) makes estimates about the child's engagement level and motivation, labeled with a therapist's resulting action, to train a robot to become progressively more autonomous when responding to the child.

Another neural network architecture used for HRI is long short-term memory (LSTM) recurrent neural networks. Unlike MLPs and other feed-forward neural networks, recurrent neural networks leverage a feedback loop that retains and uses information about previous input when processing future input. They can thus extract temporal features, which is especially important when learning over continuous data, as in HRI applications. LSTM networks, composed of LSTM cells, are able to learn long-term dependencies throughout the data stream by implementing an input gate and output gate to protect stored memory from irrelevant input (123). This is beneficial in HRI because a person's behavior may be influenced by previous interactions (e.g., a person might perform better on a task after the robot gives encouragement). For example, Dermouche & Pelachaud (111) designed an interaction loop LSTM model that takes as input the behavior of both the person and robot to continuously adapt to a user.

Neural network and deep learning approaches have proven successful in multiple areas, but the extensive amount of training data required may make them infeasible for learning the behavior of a specific user. Additionally, deep neural networks can be very sensitive to the values of hyperparameters, and care must be taken to avoid overfitting when tuning.

7. CHALLENGES AND OPPORTUNITIES FOR BEHAVIOR ADAPTATION

There are several opportunities to improve behavior adaptation to advance the efficacy of robots for longitudinal cognitive neurorehabilitation and other notable applications.

7.1. Combining Domain Knowledge with User Adaptation Approaches

There are many algorithms for learning and adapting to individual preferences, and some that enable a robot to learn desired behavior from a domain expert. Learning from experts such as a therapist or a person receiving training can help robots better serve users and ensure their behavior is well aligned with training goals. Both are important capabilities for robots, but how to best integrate these approaches is still an open problem. There has been some work to enable stakeholders—particularly those with low technology literacy, such as clinicians—to reprogram and retask robots to integrate this knowledge (see 19), but there are many opportunities to extend the scalability and automation of such approaches.

The ability to learn about and adapt to a user is essential to maximize adherence and engagement with training. Expanding the means through which robots can learn how they should modify their behavior may ultimately improve their efficacy and acceptability and increase their functionality in numerous other applications.

7.2. Imbuing Contextual Understanding

As a person continues training, they and a robot may experience many new situations. However, the majority of existing algorithms for behavior adaptation learn state–action pairs without an understanding of why those actions are appropriate. Thus, it is difficult for a robot to utilize existing knowledge in new situations, even if they are similar. Computationally, it is inefficient for a robot to have to learn how to handle every new situation and context.

Advances in transfer learning can help, but to our knowledge, little work has explored this avenue of research. There are many possibilities for providing robots with a deeper contextual understanding, thus enabling them to apply existing knowledge to new situations in order to more effectively engage with a user (124, 125).

7.3. Interacting with Groups

Another challenge is enabling robots to interact with groups of people. Many of the approaches discussed in Section 6 focus on dyadic interactions between one human and one robot. However, many situations may require robots to interact with multiple people (e.g., facilitating a group activity in nursing homes or interacting with caregivers). Interacting with groups may also help robots understand how to translate shared preferences (e.g., cultural preferences) to individual interactions. Developing new algorithms to enable adaptation to groups would improve robot utility and acceptance in many settings.

7.4. Ending Longitudinal Interactions

While the goal of behavior adaptation is often to sustain engagement throughout training, the goal of this training is often to teach people how to transfer the skills they practiced to their everyday lives. Ideally, they will no longer need a robot at the end of training, so researchers have begun to consider how a robot can alter its behavior throughout an interaction to help users prepare for its absence (46, 126). This can take many forms depending on the context, such as using a robot to facilitate interactions between a child with autism spectrum disorder and another person (e.g., a therapist or classmate) to encourage interactions without the robot (126). However, this is still an open problem in HRI in general, and particular care must be taken in the context of neurorehabilitation when working with people with cognitive impairments.

7.5. Ethical Considerations

Finally, there are many ethical considerations that must be thoroughly explored in this space. For instance, what role should a robot play in the relationship between a person receiving training and their caregivers, clinicians, and so on (e.g., companion or point of connection), and how can robot developers help facilitate the specification and modification of these roles? Throughout its use, a person may begin to see a robot as a friend and companion, so it is important to avoid overreliance on the robot so that its removal does not have detrimental effects on the person. Research shows that people can become highly emotionally attached to robots, sometimes seeing them as team members or pets; people name and dress up robotic vacuums (e.g., Roombas) in their homes (127, 128), and soldiers mourn the “deaths” of explosive ordinance disposal robots (129).

Additionally, humanlike robots have the potential to cause Turing deceptions, when a person does not know whether they are interacting with a human or a robot, particularly among people with cognitive impairments (3, 12, 130). This can be confusing for the person and may also inadvertently give a robot more authority if it resembles a caregiver or clinician. Thus, roboticists must consider how users may perceive a robot.

It is also important to consider how a robot can provide critical information to a person's caregivers and clinicians while respecting privacy. Recent advances in explainable AI can improve transparency when communicating to stakeholders why a robot behaved a certain way (see 86), but to what extent personal information should be shared is still an open problem. As robots become more ubiquitous in neurorehabilitation and other health applications, researchers must carefully consider questions such as these to avoid unintended consequences for potentially vulnerable populations (131).

8. OUR WORK TO DATE

Our team has several projects that address some of the aforementioned challenges. In particular, we have begun to address leveraging domain knowledge in robot personalization (19), learning preferences and abilities independent of specific tasks (30), and learning shared preferences from multiple users (116, 132).

8.1. Leveraging Domain Knowledge in Robot Personalization

Enabling all stakeholders, including those with low technology literacy, to reprogram robots is critical for modifying robot behavior, particularly for applications such as neurorehabilitation, where it is crucial to give control to those who may not have programming experience. Existing frameworks to support novice programmers are almost entirely procedural, require understanding code structure, and do not allow high-level specification of desired behavior.

Our system JESSIE (Just Express Specifications, Synthesize, and Interact) (see **Figure 4**) allows nonprogrammers to quickly and easily specify complex robot behavior. Thus, clinicians can specify custom behavior for personalized cognitive training regimens and reactions to keep people engaged and focused on overarching goals, rather than concerning themselves with specific implementation details or robot actions (19).

The contributions of this work were as follows. First, we presented JESSIE, which couples control synthesis methods with an accessible tangible specification interface. JESSIE enables users to specify and synthesize social robot controllers that afford personalized activities, reactions, and behavioral constraints. Second, we demonstrated JESSIE in the context of enabling clinicians to develop cognitive training regimens for people with MCI, promoting quick and easy customization.




Figure 4

The JESSIE (Just Express Specifications, Synthesize, and Interact) system. JESSIE employs control synthesis with a tangible front end to enable people to create customizable programs for social robots within the context of neurorehabilitation.

Finally, we made JESSIE open source as an artifact to support reproducibility for other robotics contexts (19).

By making the benefits of control synthesis accessible, JESSIE enabled clinicians who had no prior experience programming robots to integrate their expert knowledge into cognitive therapy sessions with personalized activities, reactions, and constraints. Our observations suggest that JESSIE enables novice programmers to leverage control synthesis techniques to create complex, interactive sessions with a social robot that would take more time to write and test with procedural programming languages. Thus, this system will enable the robotics community to customize social robot behavior, adapt to end-user preferences, and promote longitudinal HRI in numerous application domains, extending the scalability, accessibility, and personalization of social robots (19).

In the future, we will explore how autonomous behavior adaptation methods can incorporate expert knowledge, particularly for cognitive training. This will help avoid overreliance on already overburdened caregivers and extend the scalability of this system.

8.2. Task-Independent Understanding of Preferences

Enabling robots to infer task-dependent goals and preferences of users will enable better collaboration with humans and faster learning on unseen tasks, particularly when working with people who may require both physical and cognitive support from robots. To personalize behavior to the individual needs of users, robots need to learn users' unique preferences through interaction. Current preference-learning techniques lack the ability to infer long-term, task-independent preferences in interactive, incomplete-information settings.

Our preference-inference formulation, ORIRL, enables robots to infer user preferences solely by observing user behavior in various tasks. The robot's goal is to infer the user's preferences in a task-independent manner and to understand how these preferences lead to the observed behavior.

The contributions of this work were as follows. First, we presented ORIRL, which learns user preferences through observation of a user. Second, we presented an algorithm based on maximum-margin methods for performing this inference. Finally, we validated ORIRL in a realistic, long-term robot-assisted interaction study (30).

Because it relies exclusively on observational data gathered as users complete tasks, ORIRL is highly suitable for assistive robots working with users who are not robotics experts. Our novel formulation successfully infers a user's task-independent preferences and predicts features of a user's actions for unseen tasks, facilitating personalized workflows for each user. The ability to model a user's preferences across different situations over long time periods will improve personalization and collaboration between users and robots (30).

In the future, we will seek to improve the features that map each user action in order to improve system performance, learn preferences more quickly, and require less data.

8.3. Learning Shared Preferences from Multiple Users

In many scenarios, multiple agents may complete tasks in similar environments. In neurorehabilitation contexts, PwD may have similar preferences and may therefore exhibit similar reactions. By aggregating the data they collect, each agent may learn to perform their respective tasks faster by leveraging information gathered from the other agents.

We generalized the multiarmed bandit problem and formulated the ϵ -multiplayer multiarmed bandit problem, which models heterogeneous multitask learning in a multiagent setting. In this problem, a set of M players (e.g., robots) are deployed to perform similar tasks, simultaneously interacting with a set of actions/arms (e.g., an activity for PwD) and receiving feedback from different reward distributions (e.g., a PwD's personal preferences) for taking the same action;

$\epsilon \geq 0$ is a discrepancy parameter that upper bounds the pairwise distances between different reward distributions for different players on the same arm. The players can communicate and share information among one another, with a goal of minimizing their collective regret (116, 132).

The contributions of this work were threefold. First, we proposed an upper confidence bound (UCB)-based algorithm that adaptively aggregates rewards collected by different players and is robust against negative transfer. Second, we provided performance guarantees by showing that when ϵ is small, we improve our collected regret bound by a factor of M . Our algorithm also exhibits robustness; we showed a fallback guarantee that when ϵ is large and it is unsafe for the players to aggregate data aggressively, our algorithm still has a performance no worse than that of the baseline algorithm (UCB-1) (133) by a constant factor. Finally, we validated our algorithm empirically on synthetic data (116, 132).

To our knowledge, this is the first algorithm for multiplayer bandit learning that is adaptive and robust against dissimilarities between sources of data. The performance guarantees ensure that robots can leverage information acquired from multiple users without negatively impacting individuals. Thus, this approach is suitable for working with PwD, who may not have similar preferences or reactions to robot behavior (116, 132).

In the future, we are interested in exploring this approach with other learning frameworks, such as contextual bandits and MDPs. We will also evaluate our algorithms in real-world applications, including cognitive training.

9. CONCLUSION

People with cognitive impairments are in a unique position where their needs and preferences may change dramatically over the course of a training regimen. However, existing approaches either assume that a person's preferences stay constant throughout an interaction or do not take their preferences into account at all. Thus, these approaches are not necessarily appropriate when working with people with cognitive impairments. The development of new methods that consider a person's dynamic state can help improve the efficacy of robot-assisted neurorehabilitation, for dementia and beyond.

The robots and methods discussed in this review can improve existing cognitive training practices, particularly in longitudinal home settings. By building on these approaches, behavior adaptation methods can enable more engaging interactions between people and robots. Through studies with stakeholders such as PwD and their clinicians and caregivers, robots can improve engagement and sustainment to benefit people in countless contexts, from improving adherence to training regimens to bettering their daily life.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

Research reported in this paper is supported by the National Science Foundation under grants IIS-1915734 and CMMI-1935500.

LITERATURE CITED

1. World Health Organ. 2011. *World report on disability 2011*. Rep., World Health Organ., Geneva
2. Riek LD. 2017. Healthcare robotics. *Commun. ACM* 60(11):68–78

3. Moharana S, Panduro AE, Lee HR, Riek LD. 2019. Robots for joy, robots for sorrow: community based robot design for dementia caregivers. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 458–67. Piscataway, NJ: IEEE
4. Natl. Alliance Caregiv., AARP Public Policy Inst. 2015. *Caregiving in the U.S. Rep.*, Natl. Alliance Caregiv. and AARP Public Policy Inst., Washington, DC
5. World Health Organ. 2015. *World report on ageing and health*. Rep., World Health Organ., Geneva
6. Wada K, Shibata T, Musha T, Kimura S. 2008. Robot therapy for elders affected by dementia. *IEEE Eng. Med. Biol. Mag.* 27(4):53–60
7. Krebs H, Volpe B, Aisen M, Hogan N. 2000. Increasing productivity and quality of care: robot-aided neuro-rehabilitation. *J. Rehabil. Res. Dev.* 37:639–52
8. Semprini M, Laffranchi M, Sanguineti V, Avanzino L, De Icco R, et al. 2018. Technological approaches for neurorehabilitation: from robotic devices to brain stimulation and beyond. *Front. Neurol.* 9:212
9. Donati AR, Shokur S, Morya E, Campos DS, Moiola RC, et al. 2016. Long-term training with a brain-machine interface-based gait protocol induces partial neurological recovery in paraplegic patients. *Sci. Rep.* 6:30383
10. Pineau J, Montemerlo M, Pollack M, Roy N, Thrun S. 2003. Towards robotic assistants in nursing homes: challenges and results. *Robot. Auton. Syst.* 42:271–81
11. Pino O, Palestra G, Trevino R, De Carolis B. 2020. The humanoid robot NAO as trainer in a memory program for elderly people with mild cognitive impairment. *Int. J. Soc. Robot.* 12:21–33
12. Riek LD. 2016. Robotics technology in mental health care. In *Artificial Intelligence in Behavioral and Mental Health Care*, ed. DD Luxton, pp. 185–203. London: Academic
13. Stefano M, Patrizia P, Mario A, Ferlini G, Rizzello R, Rosati G. 2014. Robotic upper limb rehabilitation after acute stroke by NeReBot: evaluation of treatment costs. *BioMed Res. Int.* 2014:265634
14. Guan C, Bouzida A, Oncy-Avila R, Moharana S, Riek LD. 2021. Taking an (embodied) cue from community health: designing dementia caregiver support technology to advance health equity. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pap. 655. New York: ACM
15. Țăpuș A, Țăpuș C, Matarić MJ. 2009. The use of socially assistive robots in the design of intelligent cognitive therapies for people with dementia. In *2009 IEEE International Conference on Rehabilitation Robotics*, pp. 924–29. Piscataway, NJ: IEEE
16. Graf B, Hans M, Schraft RD. 2004. Care-O-bot II—development of a next generation robotic home assistant. *Auton. Robots* 16:193–205
17. Kittmann R, Fröhlich T, Schäfer J, Reiser U, Weißhardt F, Haug A. 2015. Let me introduce myself: I am Care-O-bot 4, a gentleman-bot. In *Mensch und Computer 2015 Tagungsband*, ed. S Diefenbach, N Henze, M Pielot, pp. 223–32. Stuttgart, Ger.: Oldenbourg
18. Asprino L, Gangemi A, Nuzzolese AG, Presutti V, Recupero DR, Russo A. 2019. Ontology-based knowledge management for comprehensive geriatric assessment and reminiscence therapy on social robots. In *Data Science for Healthcare: Methodologies and Applications*, ed. S Consoli, DR Recupero, M Petković, pp. 173–93. Cham, Switz.: Springer
19. Kubota A, Peterson EI, Rajendren V, Kress-Gazit H, Riek LD. 2020. JESSIE: synthesizing social robot behaviors for personalized neurorehabilitation and beyond. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 121–30. New York: ACM
20. Clabaugh C, Matarić MJ. 2018. Robots for the people, by the people: personalizing human-machine interaction. *Sci. Robot.* 3:eaat7451
21. Moyle W, Jones CJ, Murfield JE, Thalib L, Beattie ER, et al. 2017. Use of a robotic seal as a therapeutic tool to improve dementia symptoms: a cluster-randomized controlled trial. *J. Am. Med. Dir. Assoc.* 18:766–73
22. Geva N, Uzefovsky F, Levy-Tzedek S. 2020. Touching the social robot PARO reduces pain perception and salivary oxytocin levels. *Sci. Rep.* 10:9814
23. Țăpuș A, Țăpuș C, Matarić M. 2009. The role of physical embodiment of a therapist robot for individuals with cognitive impairments. In *RO-MAN 2009: The 18th IEEE International Symposium on Robot and Human Interactive Communication*, pp. 103–7. Piscataway, NJ: IEEE

24. Cruz-Sandoval D, Morales-Tellez A, Sandoval EB, Favela J. 2020. A social robot as therapy facilitator in interventions to deal with dementia-related behavioral symptoms. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 161–69. New York: ACM
25. Barman A, Chatterjee A, Bhide R. 2016. Cognitive impairment and rehabilitation strategies after traumatic brain injury. *Indian J. Psychol. Med.* 38:172–81
26. Huckans M, Hutson L, Twamley E, Jak A, Kaye J, Storzbach D. 2013. Efficacy of cognitive rehabilitation therapies for mild cognitive impairment (MCI) in older adults: working toward a theoretical model and evidence-based interventions. *Neuropsychol. Rev.* 23:63–80
27. Lopresti EF, Mihailidis A, Kirsch N. 2004. Assistive technology for cognitive rehabilitation: state of the art. *Neuropsychol. Rehabil.* 14:5–39
28. Chan J, Nejat G. 2012. Social intelligence for a robot engaging people in cognitive training activities. *Int. J. Adv. Robot. Syst.* 9:113
29. Tsiakas K, Abujelala M, Makedon F. 2018. Task engagement as personalization feedback for socially-assistive robots and cognitive training. *Technologies* 6:49
30. Woodworth B, Ferrari F, Zosa TE, Rick LD. 2018. Preference learning in assistive robotics: observational repeated inverse reinforcement learning. In *Proceedings of the 3rd Machine Learning for Healthcare Conference*, ed. F Doshi-Velez, J Fackler, K Jung, D Kale, R Ranganath, et al., pp. 420–39. Proc. Mach. Learn. Res. 85. N.p.: PMLR
31. Rossi S, Ferland F, Tăpus A. 2017. User profiling and behavioral adaptation for HRI: a survey. *Pattern Recognit. Lett.* 99:3–12
32. Wang L, Rau PLP, Evers V, Robinson BK, Hinds P. 2010. When in Rome: the role of culture & context in adherence to robot recommendations. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 359–66. Piscataway, NJ: IEEE
33. Tăpus A, Țăpuș C, Matarić MJ. 2008. User—robot personality matching and assistive robot behavior adaptation for post-stroke rehabilitation therapy. *Intell. Serv. Robot.* 1:169
34. World Health Organ. 2017. *Global action plan on the public health response to dementia 2017–25*. Rep., World Health Organ., Geneva
35. Teng E, Tassniyom K, Lu PH. 2012. Reduced quality-of-life ratings in mild cognitive impairment: analyses of subject and informant responses. *Am. J. Geriatr. Psychiatry* 20:1016–25
36. Petersen RC. 2004. Mild cognitive impairment as a diagnostic entity. *J. Intern. Med.* 256:183–94
37. Bruscoli M, Lovestone S. 2004. Is MCI really just early dementia? A systematic review of conversion studies. *Int. Psychogeriatr.* 16:129–40
38. Arnáiz E, Almkvist O. 2003. Neuropsychological features of mild cognitive impairment and preclinical Alzheimer's disease. *Acta Neurol. Scand.* 107:34–41
39. Fleming JM, Shum D, Strong J, Lightbody S. 2005. Prospective memory rehabilitation for adults with traumatic brain injury: a compensatory training programme. *Brain Injury* 19:1–10
40. Bahar-Fuchs A, Clare L, Woods B. 2013. Cognitive training and cognitive rehabilitation for persons with mild to moderate dementia of the Alzheimer's or vascular type: a review. *Alzheimer's Res. Ther.* 5:35
41. Garand L, Dew MA, Eazor LR, DeKosky ST, Reynolds CF III. 2005. Caregiving burden and psychiatric morbidity in spouses of persons with mild cognitive impairment. *Int. J. Geriatr. Psychiatry* 20:512–22
42. Dixon E, Lazar A. 2020. Approach matters: linking practitioner approaches to technology design for people with dementia. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York: ACM. <https://doi.org/10.1145/3313831.3376432>
43. Poulos CJ, Bayer A, Beaupre L, Clare L, Poulos RG, et al. 2017. A comprehensive approach to reablement in dementia. *Alzheimer's Dement. Transl. Res. Clin. Interv.* 3:450–58
44. Clare L, van Paasschen J, Evans SJ, Parkinson C, Woods RT, Linden DE. 2009. Goal-oriented cognitive rehabilitation for an individual with mild cognitive impairment: behavioural and neuroimaging outcomes. *Neurocase* 15:318–31
45. Lehmann H, Iacono I, Robins B, Marti P, Dautenhahn K. 2011. 'Make it move': playing cause and effect games with a robot companion for children with cognitive disabilities. In *Proceedings of the 29th Annual European Conference on Cognitive Ergonomics*, pp. 105–12. New York: ACM

46. Robins B, Dautenhahn K, Te Boekhorst R, Billard A. 2005. Robotic assistants in therapy and education of children with autism: Can a small humanoid robot help encourage social interaction skills? *Univers. Access Inf. Soc.* 4:105–20
47. Scassellati B. 2007. How social robots will help us to diagnose, treat, and understand autism. In *Robotics Research: Results of the 12th International Symposium ISRR*, ed. S Thrun, R Brooks, H Durrant-Whyte, pp. 552–63. Berlin: Springer
48. Kim ES, Berkovits LD, Bernier EP, Leyzberg D, Shic F, et al. 2013. Social robots as embedded reinforcers of social behavior in children with autism. *J. Autism Dev. Disord.* 43:1038–49
49. Cabibihan JJ, Javed H, Ang M, Aljunied SM. 2013. Why robots? A survey on the roles and benefits of social robots in the therapy of children with autism. *Int. J. Soc. Robot.* 5:593–618
50. Ujike S, Yasuhara Y, Osaka K, Sato M, Catangui E, et al. 2019. Encounter of Pepper-CPGE for the elderly and patients with schizophrenia: an innovative strategy to improve patient's recreation, rehabilitation, and communication. *J. Med. Investig.* 66:50–53
51. Raffard S, Bortolon C, Khoramshahi M, Salesse RN, Burca M, et al. 2016. Humanoid robots versus humans: How is emotional valence of facial expressions recognized by individuals with schizophrenia? An exploratory study. *Schizophr. Res.* 176:506–13
52. Giannopulu I. 2013. Multimodal cognitive nonverbal and verbal interactions: the neurorehabilitation of autistic children via mobile toy robots. *LARIA Int. J. Adv. Life Sci.* 5:214–22
53. Costa S, Santos C, Soares F, Ferreira M, Moreira F. 2010. Promoting interaction amongst autistic adolescents using robots. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, pp. 3856–59. Piscataway, NJ: IEEE
54. Eriksson J, Matarić MJ, Winstein CJ. 2005. Hands-off assistive robotics for post-stroke arm rehabilitation. In *9th International Conference on Rehabilitation Robotics, 2005*, pp. 21–24. Piscataway, NJ: IEEE
55. Deng E, Mutlu B, Matarić MJ. 2019. Embodiment in socially interactive robots. *Found. Trends Robot.* 7:251–356
56. Kidd CD, Breazeal C. 2008. Robots at home: understanding long-term human-robot interaction. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3230–35. Piscataway, NJ: IEEE
57. Moyle W, Jones C, Cooke M, O'Dwyer S, Sung B, Drummond S. 2014. Connecting the person with dementia and family: a feasibility study of a telepresence robot. *BMC Geriatr.* 14:7
58. Gross HM, Schroeter C, Mueller S, Volkhardt M, Einhorn E, et al. 2011. Progress in developing a socially assistive mobile home robot companion for the elderly with mild cognitive impairment. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2430–37. Piscataway, NJ: IEEE
59. Tamura T, Yonemitsu S, Itoh A, Oikawa D, Kawakami A, et al. 2004. Is an entertainment robot useful in the care of elderly people with severe dementia? *J. Gerontol. A* 59:M83–85
60. Yamazaki R, Kochi M, Zhu W, Kase H. 2018. A pilot study of robot reminiscence in dementia care. *Int. J. Med. Health Biomed. Bioeng. Pharm. Eng.* 12:253–57
61. Libin A, Cohen-Mansfield J. 2004. Therapeutic robot for nursing home residents with dementia: preliminary inquiry. *Am. J. Alzheimer's Dis. Other Dement.* 19:111–16
62. Chu MT, Khosla R, Khaksar SMS, Nguyen K. 2017. Service innovation through social robot engagement to improve dementia care quality. *Assist. Technol.* 29:8–18
63. Tao C, Han R, Huang J, Wang X, Ma L. 2015. Development and experiment study of an intelligent walking-aid robot. *Int. J. Model. Identif. Control* 24:216–23
64. Chang WL, Šabanovic S, Huber L. 2013. Use of seal-like robot PARO in sensory group therapy for older adults with dementia. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 101–2. Piscataway, NJ: IEEE
65. Šabanović S, Bennett CC, Chang WL, Huber L. 2013. PARO robot affects diverse interaction modalities in group sensory therapy for older adults with dementia. In *2013 IEEE 13th International Conference on Rehabilitation Robotics*. Piscataway, NJ: IEEE. <https://doi.org/10.1109/ICORR.2013.6650427>
66. Woods B, O'Philbin L, Farrell EM, Spector AE, Orrell M. 2018. Reminiscence therapy for dementia. *Cochrane Database Syst. Rev.* 3:CD001120
67. Wobbrock JO, Kane SK, Gajos KZ, Harada S, Froehlich J. 2011. Ability-based design: concept, principles and examples. *ACM Trans. Access. Comput.* 3:9

68. Morris MR. 2020. AI and accessibility. *Commun. ACM* 63(6):35–37
69. Whittaker M, Alper M, Bennett CL, Hendren S, Kaziunas L, et al. 2019. *Disability, bias, and AI*. Rep., AI Now Inst., N.Y. Univ., New York
70. Holthe T, Halvorsrud L, Karterud D, Hoel KA, Lund A. 2018. Usability and acceptability of technology for community-dwelling older adults with mild cognitive impairment and dementia: a systematic literature review. *Clin. Interv. Aging* 13:863–86
71. Small GW, Rabins PV, Barry PP, Buckholtz NS, DeKosky ST, et al. 1997. Diagnosis and treatment of Alzheimer disease and related disorders: consensus statement of the American Association for Geriatric Psychiatry, the Alzheimer's Association, and the American Geriatrics Society. *JAMA* 278:1363–71
72. Rocca WA, Boyd CM, Grossardt BR, Bobo WV, Rutten LJF, et al. 2014. Prevalence of multimorbidity in a geographically defined American population: patterns by age, sex, and race/ethnicity. *Mayo Clin. Proc.* 89:1336–49
73. Heerink M. 2011. How elderly users of a socially interactive robot experience adaptiveness, adaptability and user control. In *2011 IEEE 12th International Symposium on Computational Intelligence and Informatics (CINTI)*, pp. 79–84. Piscataway, NJ: IEEE
74. Martins GS, Al Tair H, Santos L, Dias J. 2019. α POMDP: POMDP-based user-adaptive decision-making for social robots. *Pattern Recognit. Lett.* 118:94–103
75. Charlton JI. 2000. *Nothing About Us Without Us: Disability Oppression and Empowerment*. Berkeley: Univ. Calif. Press
76. Lee HR, Riek LD. 2018. Reframing assistive robots to promote successful aging. *ACM Trans. Human-Robot Interact.* 7:11
77. Van Wynsberghe A. 2013. Designing robots for care: care centered value-sensitive design. *Sci. Eng. Ethics* 19:407–33
78. Castellano G, Aylett R, Dautenhahn K, Paiva A, McOwan PW, Ho S. 2008. *Long-term affect sensitive and socially interactive companions*. Paper presented at the 4th International Workshop on Human-Computer Conversation, Oxford, UK, Oct. 6–7
79. Nahum-Shani I, Hekler EB, Spruijt-Metz D. 2015. Building health behavior models to guide the development of just-in-time adaptive interventions: a pragmatic framework. *Health Psychol.* 34(Suppl.):1209–19
80. Fong T, Nourbakhsh I, Dautenhahn K. 2003. A survey of socially interactive robots. *Robot. Auton. Syst.* 42:143–66
81. Szafrir D, Mutlu B. 2012. Pay attention! Designing adaptive agents that monitor and improve user engagement. In *CHI '12: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 11–20. New York: ACM
82. Duffy BR. 2003. Anthropomorphism and the social robot. *Robot. Auton. Syst.* 42:177–90
83. Martin JN, Nakayama TK. 2013. *Intercultural Communication in Contexts*. New York: McGraw-Hill. 6th ed.
84. Foukarakis M, Leonidis A, Antona M, Stephanidis C. 2014. Combining finite state machine and decision-making tools for adaptable robot behavior. In *International Conference on Universal Access in Human-Computer Interaction: Aging and Assistive Environments*, ed. C Stephanidis, M Antona, pp. 625–35. Cham, Switz.: Springer
85. Meng Q, Wu W. 2008. Artificial emotional model based on finite state machine. *J. Central South Univ. Technol.* 15:694–99
86. Lee MH, Siewiorek DP, Smailagic A, Bernardino A, Bermúdez i Badia S. 2020. An exploratory study on techniques for quantitative assessment of stroke rehabilitation exercises. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, pp. 303–7. New York: ACM
87. Liu C, Conn K, Sarkar N, Stone W. 2008. Online affect detection and robot behavior adaptation for intervention of children with autism. *IEEE Trans. Robot.* 24:883–96
88. Ritschel H, André E. 2017. Real-time robot personality adaptation based on reinforcement learning and social signals. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 265–66. New York: ACM
89. Lockerd A, Breazeal C. 2004. Tutelage and socially guided robot learning. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vol. 4, pp. 3475–80. Piscataway, NJ: IEEE

90. Prommer T, Holzapfel H, Waibel A. 2006. Rapid simulation-driven reinforcement learning of multi-modal dialog strategies in human-robot interaction. In *INTERSPEECH 2006: Ninth International Conference on Spoken Language Processing*, pp. 1918–21. Baixas, Fr.: Int. Speech Commun. Assoc.
91. Krsmanovic F, Spencer C, Jurafsky D, Ng AY. 2006. Have we met? MDP based speaker ID for robot dialogue. In *INTERSPEECH 2006: Ninth International Conference on Spoken Language Processing*, pp. 461–64. Baixas, Fr.: Int. Speech Commun. Assoc.
92. Malfaz M, Castro-González Á, Barber R, Salichs MA. 2011. A biologically inspired architecture for an autonomous and social robot. *IEEE Trans. Auton. Ment. Dev.* 3:232–46
93. Ferreira E, Lefevre F. 2015. Reinforcement-learning based dialogue system for human-robot interactions with socially-inspired rewards. *Comput. Speech Lang.* 34:256–74
94. Karami AB, Sehaba K, Encelle B. 2016. Adaptive artificial companions learning from users' feedback. *Adapt. Behav.* 24:69–86
95. Chen H, Park HW, Breazeal C. 2020. Teaching and learning with children: impact of reciprocal peer learning with a social robot on children's learning and emotive engagement. *Comput. Educ.* 150:103836
96. Park HW, Grover I, Spaulding S, Gomez L, Breazeal C. 2019. A model-free affective reinforcement learning approach to personalization of an autonomous social robot companion for early literacy education. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, pp. 687–94. Palo Alto, CA: AAAI Press
97. Gao AY, Barendregt W, Castellano G. 2017. *Personalised human-robot co-adaptation in instructional settings using reinforcement learning*. Paper presented at the Workshop on Persuasive Embodied Agents for Behavior Change, International Conference on Intelligent Virtual Agents, Stockholm, Aug. 27
98. Biancardi B, Mancini M, Lerner P, Pelachaud C. 2019. Managing an agent's self-presentational strategies during an interaction. *Front. Robot. AI* 6:93
99. Gordon G, Spaulding S, Westlund JK, Lee JJ, Plummer L, et al. 2016. Affective personalization of a social robot tutor for children's second language skills. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pp. 3951–57. Palo Alto, CA: AAAI Press
100. Taha T, Miró JV, Dissanayake G. 2011. A POMDP framework for modelling human interaction with assistive robots. In *2011 IEEE International Conference on Robotics and Automation*, pp. 544–49. Piscataway, NJ: IEEE
101. Karami AB, Jeanpierre L, Mouaddib AI. 2009. Partially observable Markov decision process for managing robot collaboration with human. In *2009 21st IEEE International Conference on Tools with Artificial Intelligence*, pp. 518–21. Piscataway, NJ: IEEE
102. Spaulding S, Breazeal C. 2017. *Learning behavior policies for interactive educational play*. Paper presented at the Workshop on Mathematical Models, Algorithms, and Human-Robot Interaction, Robotics: Science and Systems XIII, Cambridge, MA, July 12–16
103. Hoey J, von Bertoldi A, Poupart P, Mihailidis A. 2007. Assisting persons with dementia during hand-washing using a partially observable Markov decision process. In *Proceedings of the 5th International Conference on Computer Vision Systems*. Bielefeld, Ger.: Appl. Comput. Sci. Group, Bielefeld Univ. <https://doi.org/10.2390/biecoll-icvs2007-89>
104. Hemminahaus J, Kopp S. 2017. Towards adaptive social behavior generation for assistive robots using reinforcement learning. In *2017 12th ACM/IEEE International Conference on Human-Robot Interaction*, pp. 332–40. Piscataway, NJ: IEEE
105. Belpaeme T, Baxter P, Read R, Wood R, Cuayáhuil H, et al. 2013. Multimodal child-robot interaction: building social bonds. *J. Hum.-Robot Int.* 1:33–53
106. Mitsunaga N, Smith C, Kanda T, Ishiguro H, Hagita N. 2006. Robot behavior adaptation for human-robot interaction based on policy gradient reinforcement learning. *J. Robot. Soc. Jpn.* 24:820–29
107. Chandramohan S, Geist M, Lefevre F, Pietquin O. 2011. User simulation in dialogue systems using inverse reinforcement learning. In *INTERSPEECH 2011: 12th Annual Conference of the International Speech Communication Association*, pp. 1025–28. Baixas, Fr.: Int. Speech Commun. Assoc.
108. Boularias A, Chinaei HR, Chaib-draa B. 2010. *Learning the reward model of dialogue pomdps from data*. Paper presented at the Workshop on Machine Learning for Assistive Techniques, 24th Conference on Neural Information Processing Systems, Vancouver, Can., Dec. 6–11

109. Sugiyama H, Meguro T, Minami Y. 2012. Preference-learning based inverse reinforcement learning for dialog control. In *INTERSPEECH 2012: 13th Annual Conference of the International Speech Communication Association*, pp. 222–25. Baixas, Fr.: Int. Speech Commun. Assoc.
110. Senft E, Baxter P, Kennedy J, Belpaeme T. 2015. SPARC: Supervised Progressively Autonomous Robot Competencies. In *International Conference on Social Robotics*, ed. A Tapus, E André, JC Martin, F Ferland, M Ammi, pp. 603–12. Cham, Switz.: Springer
111. Dermouche S, Pelachaud C. 2019. Generative model of agent's behaviors in human-agent interaction. In *2019 International Conference on Multimodal Interaction*, pp. 375–84. New York: ACM
112. Yan H, Ang MH, Poo AN. 2014. A survey on perception methods for human–robot interaction in social robots. *Int. J. Soc. Robot.* 6:85–119
113. Pantic M, Pentland A, Nijholt A, Huang TS. 2007. Human computing and machine understanding of human behavior: a survey. In *Artificial Intelligence for Human Computing*, ed. TS Huang, A Nijholt, M Pantic, A Pentland, pp. 44–71. Berlin: Springer
114. Bohren J, Cousins S. 2010. The SMACH high-level executive. *IEEE Robot. Autom. Mag.* 17(4):18–20
115. Sutton RS, Barto AG. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press
116. Wang Z, Singh MK, Zhang C, Riek LD, Chaudhuri K. 2020. *Stochastic multi-player bandit learning from player-dependent feedback*. Paper presented at the Workshop on Real World Experiment Design and Active Learning, 37th International Conference on Machine Learning, virtual, July 18
117. Auer P, Cesa-Bianchi N, Freund Y, Schapire RE. 2002. The nonstochastic multiarmed bandit problem. *SIAM J. Comput.* 32:48–77
118. Sutton RS, McAllester DA, Singh SP, Mansour Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems 12*, ed. Solla, T Leen, KM Müller, pp. 1057–63. Cambridge, MA: MIT Press
119. Smallwood RD, Sondik EJ. 1973. The optimal control of partially observable Markov processes over a finite horizon. *Oper. Res.* 21:1071–88
120. Dietterich TG. 2000. Hierarchical reinforcement learning with the MAXQ value function decomposition. *J. Artif. Intell. Res.* 13:227–303
121. Janeczek A, Gansterer W, Demel M, Ecker G. 2008. On the relationship between feature selection and classification accuracy. In *Proceedings of the Workshop on New Challenges for Feature Selection in Data Mining and Knowledge Discovery at ECML/PKDD 2008*, ed. Y Saeys, H Liu, I Inza, L Wehenkel, Y Van de Pee, pp. 90–105. Proc. Mach. Learn. Res. 4. N.p.: PMLR
122. Gardner MW, Dorling S. 1998. Artificial neural networks (the multilayer perceptron)—a review of applications in the atmospheric sciences. *Atmos. Environ.* 32:2627–36
123. Hochreiter S, Schmidhuber J. 1997. Long short-term memory. *Neural Comput.* 9:1735–80
124. O'Connor MF, Riek LD. 2015. Detecting social context: a method for social event classification using naturalistic multimodal data. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*. Piscataway, NJ: IEEE. <https://doi.org/10.1109/FG.2015.7284843>
125. Nigam A, Riek LD. 2015. Social context perception for mobile robots. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3621–27. Piscataway, NJ: IEEE
126. Dautenhahn K, Nehaniv CL, Walters ML, Robins B, Kose-Bagci H, et al. 2009. Kaspar – a minimally expressive humanoid robot for human–robot interaction research. *Appl. Bionics Biomech.* 6:369–97
127. Sung JY, Guo L, Grinter RE, Christensen HI. 2007. “My Roomba is Rambo”: intimate home appliances. In *UbiComp 2007: Ubiquitous Computing*, ed. J Krumm, GD Abowd, A Seneviratne, T Strang, pp. 145–62. Berlin: Springer
128. Forlizzi J, DiSalvo C. 2006. Service robots in the domestic environment: a study of the Roomba vacuum in the home. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction*, pp. 258–65. New York: ACM
129. Carpenter J. 2016. *Culture and Human-Robot Interaction in Militarized Spaces: A War Story*. London: Routledge
130. Riek L, Howard D. 2014. *A code of ethics for the human-robot interaction profession*. Paper presented at We Robot, Coral Gables, FL, Apr. 4–5

131. Kubota A, Pourebadi M, Banh S, Kim S, Riek LD. 2021. *Somebody that I used to know: the risks of personalizing robots for dementia care*. Paper presented at We Robot, Coral Gables, FL, Sept. 23–25
132. Wang Z, Zhang C, Singh MK, Riek L, Chaudhuri K. 2021. Multitask bandit learning through heterogeneous feedback aggregation. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, ed. A Banerjee, K Fukumizu, pp. 1531–39. Proc. Mach. Learn. Res. 130. N.p.: PMLR
133. Auer P, Cesa-Bianchi N, Fischer P. 2002. Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.* 47:235–56