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 92093, USA; email: akubota@ucsd.edu, lriek@ucsd.edu

Xxxx. Xxx. Xxx. Xxx. 2021. AA:1–27

https://doi.org/10.1146/((please add article doi))

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 to working in this space. We then examine state-of-the-art technical approaches to 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 they continue to provide more engaging, effective, and personalized interactions between people and robots.

Figure 1

Robots used to support people with cognitive impairments vary widely in morphology, including mobile, tabletop, humanoid, mechanistic, and zoomorphic. From left to right: Bandit (1) (Provided by Maja Matarić), Care-O-bot (2) (Provided by Fraunhofer IPA (3)), KOMPAÏ-2 (4) (Provided by KOMPAÏ Robotics (5)), Kuri (6) (Provided by Mayfield Robotics), Mabu (7) (Provided by Catalia Health), PARO (8) (Provided by Carlton SooHoo (9)).

1. INTRODUCTION

15-20% of the world's population has a disability which greatly impacts their independence (10). This can affect their ability to perform activities of daily living (ADLs) (e.g. eating, bathing), and instrumental ADLs (IADLs) (e.g. managing medication, finances) (11, 12). However, the number of people who need support exceed the availability and resources of full-time care providers, and informal caregivers (e.g. family) must often assume much of the care responsibility (13, 14), yet are provided few resources to do so, leading to stress and burnout (12).

Robots have shown great potential to help people across numerous aspects of health and wellness. Examples range across many of settings, including homes, clinics, and hospitals, and different tasks, including reducing clinician workload, supporting people with disabilities, and supporting caregivers (8, 11, 12, 15, 16, 17, 18, 19, 20).

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 reduce the cost of treatment for patients, as they take less of a clinician's time (21). 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 (22).

Across both the research and commercial sector, socially assistive robots (SARs), which provide assistance through social interactions, are being deployed in people's homes to support their health and prolong their independence. **Figure 1** shows a number of 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 PwD and their caregivers (8, 23), and can also alleviate pain and improve mood (24). Researchers are also exploring SARs to help people with cognitive impairments learn to manage their condition through cognitive training (6, 19, 25, 26).

In cognitive training and other behavioral treatments, it is critical to personalize training to an individual's 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 a training (27, 28). These treatments are traditionally led by a human neuropsychologist or cognitive therapist who works closely with a person to determine their needs and goals, and tailor training to them.

Recently, researchers explored how to enable robots to help facilitate training while supporting clinicians and caregivers (6, 12, 25, 29, 30, 31, 32, 26). 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 to social interaction strategies (e.g. initiative, personality), and preferred behaviors may vary across cultural and personal backgrounds (33, 34). 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 (6).

Longitudinal behavior adaptation methods from human-robot interaction (HRI) can help address these challenges. These approaches enable robots to create a model of a person (e.g. personality, preferences) (33) to guide how it interacts and responds them. Doing this accurately and consistently is critical when working with people with neurodegenerative conditions, such as people with dementia (PwD) (6, 35).

In this review, we explore longitudinal adaptation methods which 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 neurorehabiitation. 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). In Section 5, we examine some common communication modalities, then explore technical approaches to behavior adaptation in Section 6. To conclude, we discuss our recent work on robot behavior adaptation (Section 8) and explore some open problems in this area (Section 7).

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. Thus, robots can more effectively help people accomplish a multitude of long-term goals from overcoming memory challenges to living a healthier lifestyle.

2. MILD COGNITIVE IMPAIRMENT AND NEUROREHABILITATION

2.1. Mild cognitive impairment

Dementia is an irreversible syndrome that entails noticeable decline of cognitive function (12, 36). Approximately 11% of people aged over 65 are impacted by dementia, and each case is unique. Symptoms can range across the spectrum, from early stage (e.g. forgetfulness) to late stage (e.g. difficulty recognizing friends and family). It 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 cause reduced quality of life to family members when adopting the role of informal caregivers (37).

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 (28, 38). An estimated 20% of adults aged over 65 experience MCI, approximately 10% of whom convert to some type of dementia each year (28, 39). To date, no existing pharmacological treatments have proven effective for slowing or preventing this conversion, but studies suggest that behavioral treatments can help (28).

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 (38, 40). 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 (28). However, as people lose their independence, it can severely impact their quality of life (41, 28). It can also adversely affect their family members, put strain on their relationship with the person with MCI, and cause stress (12, 42, 28, 43, 44). This change in lifestyle and role can cause feelings of guilt, anxiety, and depression in a person with MCI and their caregivers (42, 43).

2.2. Neurorehabilitation

Many researchers have explored strategies to promote the reablement of PwD, or mitigating the impact of dementia on their function to promote independence (45). In particular, non-pharmacological approaches such as behavioral therapy can slow the onset of MCI, which can prolong independence and maintain quality of life (28). 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 (CCT) and restorative cognitive training (28).

CCT teaches a person with MCI metacognitive strategies to help bypass impaired function and minimize its impact on daily life (28, 41). 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 keep and check 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. CCT 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 (28). In our work, we focus on employing CCT with a robot (6).

In contrast, restorative cognitive training 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 (46). Furthermore, these skills typically do not generalize well (i.e. transfer) to other tasks (28).

3. ROBOTS FOR NEUROREHABILITATION

Robots for physical neurorehabilitation typically help people by physically supporting or correcting movement with the goal of restoring neuromotor function, e.g., restoring or supplementing limb function in people who had a spinal cord injury or stroke (15, 16, 17). These robots take many forms, such as robotic arms to help people control their arms and hands to complete ADL tasks (15, 16), or exoskeletons to help people walk (17).

Researchers also use SARs to support cognitive neurorehabilitation.[1] 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 (1, 47).

In addition to dementia, researchers have increasingly explored the use of SARs to support people with social and developmental disorders, particularly children with autism spectrum disorder (ASD) or attention-deficit/hyperactivity disorder (ADHD) (48, 49, 50, 51) and people with schizophrenia (20, 52, 53). For instance, children with ASD 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 (54, 55). Robots can also help improve communication between older adults with schizophrenia and their medical providers, and increase their engagement with recreational activities (20, 52).

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 shown to improve attention, memory, and executive skills in people with memory impairments (27), robots have even greater potential to improve training, as their physical embodiment plays an important role in stroke patient compliance and engagement (35, 56). Robots can increase engagement and enjoyment in social interactions due to their increased capacity for richer communication as compared to virtual systems (57). 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 (58).

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 a user. For example, in the space of 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 which 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 overviews 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 MCI 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** overviews selected robots, some designed specifically for PwD and others applied to this space.

They are typically mobile robots that provide monitoring and care, helping to ease the

[1]To our knowledge, there are no commercially available robot systems to administer cognitive neurorehabilitation at the time of writing.

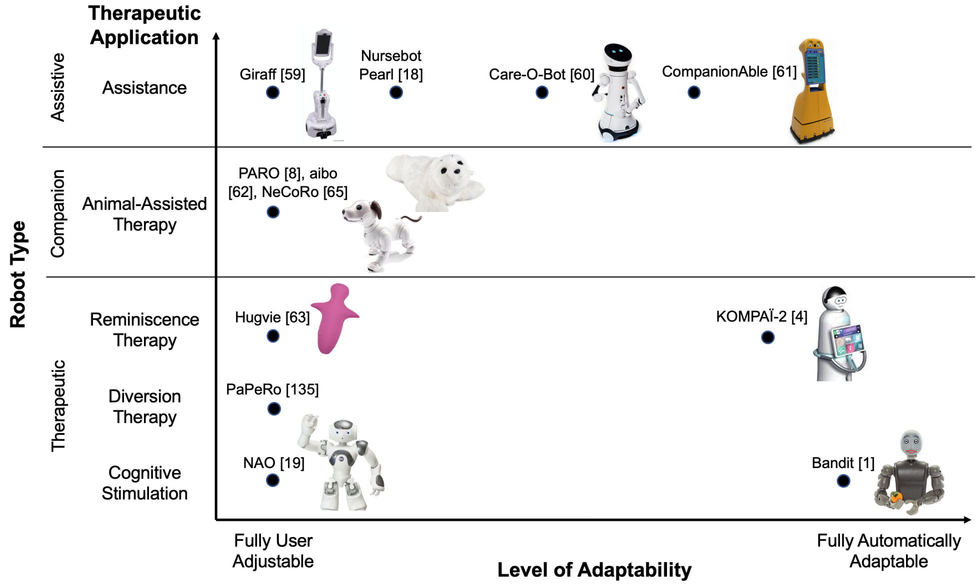

Figure 2

Exemplar robots which 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. Giraff (59) (Provided by Camanio AB), Care-O-bot (60) (Provided by Fraunhofer IPA (3)), CompanionAble (61) (Provided by Steffen Müller), PARO (8) (Provided by Carlton SooHoo (9)), aibo (62), Hugvie (63) (Provided by ATR Hiroshi Ishiguro Laboratories), KOMPAÏ-2 (4) (Provided by KOMPAÏ Robotics (5)), NAO (19), Bandit (1) (Provided by Maja Matarić).

responsibilities of informal caregivers (e.g. family, 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 (59, 61). Others include walking assistance, fetching items, or setting a table for people with mobility difficulties (60). These usually occur within a home, though some researchers are also exploring robots that can accompany users on errands outside of the home (18, 64).

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 PwD and their caregivers (8, 23). They can help stimulate interaction and serve as a point of connection between PwD and their caregivers (8). Many of these robots resemble animals, making them recognizable even to people with severe memory impairments. For instance, PARO (8) is based on a baby harp seal, and AIBO (62) 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 (8).

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 (8, 62, 65). In addition, researchers have explored using PARO to facilitate multi-sensory behavior therapy (66, 67), 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 (63, 4). Reminiscence therapy aims to help people recall long-term autobiographical memories with the aid of photographs, music, familiar objects, etc. It is highly regarded by participants and therapists, and viewed as enjoyable and effective (68). The approach is generally conversational, guided by either a human therapist or a robot itself, using a microphone and speaker in the robot to communicate with the person (63, 4). In robot-guided sessions, a robot relies on user-specific knowledge (e.g. photos from an event, 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 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 (25). Similarly, researchers have programmed humanoid robots such as the NAO for clinicians to use to assist with memory training programs (19).

4. PRINCIPLES FOR DESIGNING NEUROREHABILITATION TECHNOLOGY

When designing technology for people with disabilities, to ensure it is usable and acceptable, there are three key considerations: personalization, adaptation, and inclusion (20, 69), which are discussed below. Personalization refers to tailoring the system to an individual by considering factors such as their needs, goals, or preferences. Adaptation is the ability for 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 on potentially vulnerable populations, such as the exacerbation of disability-based bias (70, 71).

4.1. Personalization and adaptation

It is critical that neurorehabilitative technologies are personalized to an individual, from simply including their name to adapting to suit their unique preferences and abilities (72). This is important when developing any technology for people who may not be represented by a "typical" user (69). 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 (42, 73). This is likely due to the repetition and structure that defined these early interventions (e.g. memorizing and repeating a specific list of words), without a clear connection to an 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, diabetes) or have different living situations (e.g. living alone, in a nursing home) (74). If the system is not personalized to them, it 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 perfor-

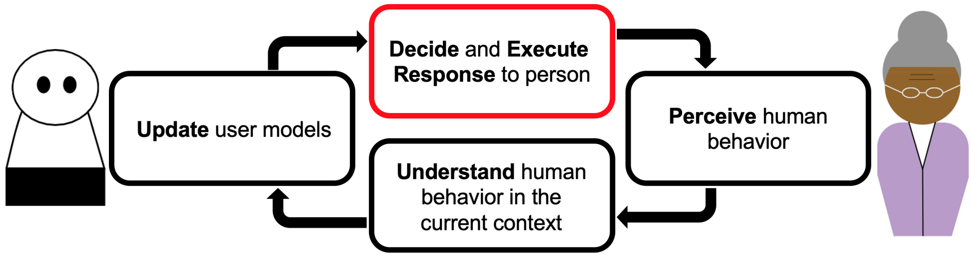

Figure 3

The general framework of systems that adapt to users. We focus on Deciding and Executing the robot's response to a person. *Figure inspired by (76).*

mance and abilities, modern neurocognitive interventions have shown to be significantly more effective and beneficial for people with cognitive impairments and their caregivers as compared to non-personalized interventions (41, 29).

The ability for technology to be personalized calls for the system to either 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.

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. Using 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.

Automatic adaptation alleviates the responsibility of continually adjusting a robot's behavior from caregivers or clinicians who can 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 (75). 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.

In order 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 that 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, observations) vs. obtain explicitly (e.g. through questionnaires, surveys)?* Section 5 overviews 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, 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, e.g., "Nothing about us without us"

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 (11, 12, 44).

Nihil de nobis, sine nobis, or "Nothing about us without us" is a prominent motto of disability activists (77). 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 ideation to testing. As roboticists oftentimes 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 (11, 78).

When co-designing 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" (45). 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 trust and maintain compassion with 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 impact informed consent (12, 20, 79). 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 efficacy of the system, particularly for health applications (80, 81). 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 (35, 82). 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.

Table 1 Robots can use a variety of sensors to perceive low-level interaction data about people, which can be used to infer high-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 understanding what a person is saying, their prosody and tone may also 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 over pressing buttons.
	Cognitive abilities	Task performance, Speech	A user's cognitive abilities can influence their goals and what treatment regimens may be most effective.

Examples of low-level behaviors that a robot may gather include their speech (e.g. what they say, how they say it), gestures and movement (e.g. human activity recognition), and physiological signals (e.g. heart or respiration rate). Performance data is typically application specific and depends on the task(s) (e.g. accuracy, 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 (33). 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).

An alternate approach is to ask a user about their preferences such as in a questionnaire or survey (33). 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 to sense about people, both low-level and higher level, for social robots for neurorehabilitation.

Once a robot perceives a person's behavior, the robot must consider how that behavior relates to a) the person's current state and/or b) 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 if the person is frustrated with or challenged by a cognitive training exercise (35).

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 (33). 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). Understanding both individual actions and translating them into a more thorough model of a person is important for personalizing robot behavior.

5.2. Synthesizing robot behavior in response to people

Effective human-robot interaction 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, role) (80, 81). In the context of neurorehabilitation, these behaviors can help improve a user's enjoyment of a training regimen and thus its efficacy (6, 31, 33).

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, an user with tremors may prefer to communicate via speech whereas someone who is non-verbal 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 (83).

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

Table 2 Robots can modify low-level behaviors to personalize high-level aspects of an interaction in order to fulfill 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 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 the anthropomorphize robots, even those that are not humanoid [31], 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, and 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. E.g., a robot may adopt a more passive communication style in countries where people tend to have more reserved communication styles [75].
	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.

regimen and improve their perceptions of the robot (31, 33). **Table 2** overviews 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.[2]

While perceiving and understanding human behavior is an important aspect 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

[2]As many social behaviors are not robot specific (e.g. dialogue), we also include select systems which were demonstrated on non-physically embodied systems, such as virtual agents.

Table 3 Common technical approaches for machines which adapt behavior to people.

Approach / Existing Work	Strengths	Limitations
Finite State Machines [58, 86, 87, 88]	• Straightforward • Existing libraries for implementation on robots (e.g. SMACH) • Good for structured, short interactions	• 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 [1, 83, 90]	• Good for reacting to continuous streams of data rather than windows of time	• 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) [19, 31, 92, 93, 94, 95, 96, 97, 98, 99, 100, 101, 102, 103]	• Model-free, or can learn a model about user behavior / preferences	• 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 [76, 109, 110, 133, 134]	• 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	• State space becomes intractable for complex interactions • Interactions generally cannot be split into discrete states
Hierarchical RL [18, 30, 111, 112]	• 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	• Does not take combinations of behaviors into consideration, so it is not guaranteed to find a globally optimum policy • Interactions generally cannot be split into discrete states
Policy Gradient RL [35, 107]	• Naturally handles continuous states and actions	• Difficult to derive an appropriate reward function (i.e. preferences from behavior) • Need to find appropriate parameter values
Inverse RL [32, 114, 115, 116]	• Learns a reward function from a human expert rather than relying on exploration of different behaviors	• Requires feedback from a human expert • Human experts do not necessarily behave optimally or rationally
Neural Networks: MLP [117]	• Hierarchical layers enable high-level feature extraction from raw or low-level input data	• 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 Networks: LSTM [118]	• 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	• Requires large amounts of training data which can make it difficult to learn from an individual person • Difficult to optimize hyperparameters

process. For a detailed survey on human behavior analysis, please refer to (84, 85).

6.1. Finite State Machines (FSM)

FSMs are a relatively straightforward approach address the behavior adaptation problem (58, 86, 87, 88). In an FSM, an interaction is broken into states which guide robot behavior. The robot transitions to the next state depending on human and environmental factors.

For instance, Kidd and Breazeal (58) used an FSM on Autom, a robotic weight loss coach. A user would engaging 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 relatively structured and short interactions, many programmers are already familiar with FSMs, and there are a number of existing libraries to implement them on robots (e.g. SMACH ("State Machine", a Python-based library) (89)). 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.

6.2. Thresholding

Another approach that roboticists use is thresholding (1, 83, 90). 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. (1) 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 (AVB)** 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, correct answers) improved, the difficulty increased, whereas it decreased if they performed poorly. The authors report 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 when engagement is low, 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 (RL)

In RL, an agent learns how to best interact with its environment to maximize its rewards (91). RL configurations are generally represented using a Markov Decision Process (MDP) defined as ($\mathcal{S}$, $\mathcal{A}$, $\mathcal{T}$, $\mathcal{R}$, γ) 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 (91). 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 to 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 (19, 31, 92, 93, 94, 95, 96, 97, 98, 99, 100, 101, 102, 103). For instance, Tsiakas et al. (31) 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 multi-armed bandit problem (101, 102, 104). The multi-armed 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 selecting behaviors in order to maximize a person's engagement, performance, etc.

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. (101) implemented the **Exponential-weight algorithm for Exploration and Exploitation (Exp3)** (105) on a Pepper robot for puzzle solving. The robot would learn the person's preference for supportive behaviors (e.g. give hints, provide 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 (31, 92). However, Q-learning and MDPs have many limitations, such as assuming the world is fully observable, and being time and storage intensive (35, 106, 107). 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 ($\mathcal{S}$, $\mathcal{A}$, $\mathcal{T}$, $\mathcal{R}$, ω, $\mathcal{O}$, γ) 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 (108). 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 (76), navigating a robotic wheelchair (109), and helping PwD wash their hands by giving visual or verbal prompts (110). 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** (18, 30, 111, 112). This approach divides the overall MDP into smaller, more manageable ones which simplifies the problem and can help reduce memory requirements (30). It can also allow for greater modularity of the system's behaviors; for instance, Chan et al. (30) used the MAXQ hierarchical RL approach (113) to abstract their system into a temporal module, state module, and subtask module which each 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 reinforcement learning (PGRL)** methods for behavior adaptation (35, 107). PGRL 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 which it evaluates according to the reward function. Then, it perturbs the policy by modifying each parameter. Finally, it evaluates the new policy, and repeats 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 which are both important aspects of behavior adaptation in HRI (107). However, researchers have reported challenges deriving an appropriate reward function to accurately translate user behavior to explicit preferences (35, 107).

A slightly different approach is **inverse reinforcement learning (IRL)** to learn how to behave from an expert agent, assumed to behave optimally (114, 115). The agent 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 rationally or optimally, and it is not always possible to discern an exact policy, so researchers have expanded IRL to help overcome these limitations (116).

Additionally, researchers have worked to infer user preferences solely from observing a user as in **Observational Repeated Inverse Reinforcement Learning (ORIRL)** (32). 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 began to leverage advances in neural networks (NN) and deep learning for robot behavior adaptation (117, 118). NNs are a broad set of algorithms inspired by biological NNs that enable agents to recognize patterns in data, generally without having to define underlying task-specific rules. NNs 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 (119).

Senft et al. (117) used a **multilayer perceptron (MLP)** to enable a robot to learn from a therapist how to interact with children with ASD. An MLP is a supervised feed-forward NN composed of multiple perceptrons, or binary classifiers, with a unique set of weights (120). The use of multiple perceptrons allows the MLP to approximate nonlinear functions for multi-class classification. The MLP used by Senft et al. estimates about the child's engagement level and motivation, labelled with a therapist's resulting action, to train a robot to become progressively more autonomous when responding to the child.

Another NN architecture used for HRI is **long short-term memory (LSTM) recurrent neural networks (RNN)**. Unlike MLPs and other feed-forward NNs, RNNs leverage a feedback loop which 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 (121). This is beneficial in HRI as 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 et al. (118) designed an Interaction Loop LSTM model which takes as input the behavior of both the person and robot to continuously adapt to a user.

NN and deep learning approaches have proven successful in a number of areas, but the extensive amount of training data required may make this approach 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.

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 is some work to enable stakeholders, particularly those with low technology literacy such as clinicians, to reprogram and re-task robots to integrate this knowledge (c.f. (6)), 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, acceptability, and further their functionality in numerous other applications.

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 explores 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 (122, 123).

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, interacting with caregivers). Interacting with groups may also help robots understand how to translate shared preferences (e.g. cultural) to individual interactions. Developing new algorithms to enable adaptation to groups would improve robot utility and acceptance in many settings.

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 (48, 124). This can take many forms depending on the context, such as using a robot to facilitate interactions between a child with ASD and another person (e.g. therapist, classmate) to encourage interactions without the robot (124). 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.

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, etc. (e.g. companion, point of connection)*, and *How can robot developers help facilitate the specifica-*

tion 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 over reliance 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 (125, 126) and soldiers mourn the “deaths” of explosive ordinance disposal (EOD) robots (127).

Additionally, humanlike robots have the potential to cause “Turing Deceptions,” or when someone does not know if they are interacting with a human or robot, particularly among people with cognitive impairments (12, 20, 128). This can confusing for a 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 artificial intelligence (AI) can improve transparency when communicating to stakeholders why a robot behaved a certain way (c.f. (88)), 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 on potentially vulnerable populations (129).

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 (6), learning preferences and abilities independent of specific tasks (32), and learning shared preferences from multiple users (104, 130).

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 non-programmers 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 (6).

The **contributions** of this work are 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 which 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. Finally, JESSIE was made open-source as an artifact to support reproducibility for other robotics contexts (6).

By making the benefits of control synthesis accessible, JESSIE enabled clinicians, who had no prior experience programming robots, to integrate their expert knowledge into cog-

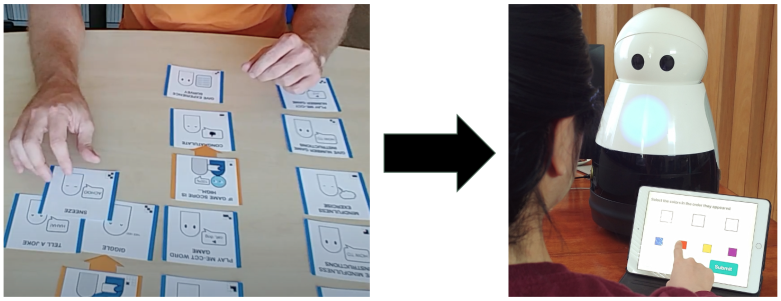

Figure 4

JESSIE employs control synthesis with a tangible front-end to enable people to create customizable programs for social robots within the context of neurorehabilitation.

nitive 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 on a social robot, which 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 (6).

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 preference 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. In order 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, Observational Repeated Inverse Reinforcement Learning (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 preference in a task-independent manner, and to understand how these preferences lead to the observed behavior.

The **contributions** of this work are 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 (32).

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 predict 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 (32).

In the future, we 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 multi-armed bandit problem and formulated the ϵ-multi-player multi-armed bandit (ϵ-MPMAB) problem which models heterogeneous multi-task learning in a multi-agent setting. In an ϵ-MPMAB problem instance, a set of M players (e.g. robots) are deployed to perform similar tasks; they simultaneously interact with a set of actions/arms (e.g. an activity for PwD), and they receive 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 (104, 130).

The **contributions** of this work are 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 show 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 (131)) by a constant factor. Finally, we validate our algorithm empirically on synthetic data (104, 130).

To our knowledge, this is the first algorithm for multi-player bandit learning that is adaptive and robust against dissimilarities between sources of data. The performance guarantees ensure robots can leverage information acquired from multiple users without negatively impacting individuals. Thus, this approach is suitable for working with PwD who we cannot necessarily assume will have similar preferences or reactions to robot behavior (104, 130).

In the future, we are interested in exploring this approach with other learning frameworks such as contextual bandits and Markov Decision Processes. 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 assume a person's preferences stay constant throughout an interaction, or do not take their preferences into account at all. Thus, they 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 Grant Nos. IIS-1915734 and CMMI-1935500.

LITERATURE CITED

1. Tapus A, Tapus C, Mataric 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. 924–929. IEEE
2. 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 robot. *Mensch und computer 2015–proceedings*
3. Fraunhofer Institute for Manufacturing Engineering and Automation. 2021. Care-o-bot 4. `https://www.care-o-bot.de/en/`, accessed 2021-03-19
4. 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*. Springer
5. Kompai Robotics. 2021. Kompai robots. `https://kompai.com/`, accessed 2021-03-19
6. 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. 121–130
7. Clabaugh C, Matarić MJ. 2018. Robots for the people, by the people: Personalizing human-machine interaction. *Science Robotics* 3:eaat7451
8. Wada K, Shibata T, Musha T, Kimura S. 2008. Robot therapy for elders affected by dementia. *IEEE Engineering in medicine and biology magazine* 27:53–60
9. Panospin360. 2021. 360° virtual tour photography. `https://www.panospin360.com/`, accessed 2021-03-19
10. World Health Organization. 2011. *World report on disability 2011*. World Health Organization
11. Riek LD. 2017. Healthcare robotics. *Communications of the ACM* 60:68–78
12. 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. 458–467. IEEE
13. for Caregiving NA. 2015. Caregiving in the us. *NAC and the AARP Public Institute. Washington DC: Greenwald & Associates*
14. Organization WH. 2015. *World report on ageing and health*. World Health Organization
15. Krebs H, Volpe B, Aisen M, Hogan N, et al. 2000. Increasing productivity and quality of care: Robot-aided neuro-rehabilitation. *Journal of rehabilitation research and development* 37:639–652
16. 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. *Frontiers in neurology* 9:212

17. Donati AR, Shokur S, Morya E, Campos DS, Moioli RC, et al. 2016. Long-term training with a brain-machine interface-based gait protocol induces partial neurological recovery in paraplegic patients. *Scientific reports* 6:30383
18. Pineau J, Montemerlo M, Pollack M, Roy N, Thrun S. 2003. Towards robotic assistants in nursing homes: Challenges and results. *Robotics and autonomous systems* 42:271–281
19. 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. *International Journal of Social Robotics* 12:21–33
20. Riek LD. 2016. Robotics technology in mental health care. In *Artificial intelligence in behavioral and mental health care*. Elsevier
21. 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 research international* 2014
22. 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. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* :1–24
23. 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. *Journal of the American Medical Directors Association* 18:766–773
24. Geva N, Uzefovsky F, Levy-Tzedek S. 2020. Touching the social robot paro reduces pain perception and salivary oxytocin levels. *Scientific Reports* 10:1–15
25. Tapus A, Tapus C, Mataric 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. 103–107. IEEE
26. 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. 161–169
27. Barman A, Chatterjee A, Bhide R. 2016. Cognitive impairment and rehabilitation strategies after traumatic brain injury. *Indian journal of psychological medicine* 38:172
28. 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. *Neuropsychology review* 23:63–80
29. Frank Lopresti E, Mihailidis A, Kirsch N. 2004. Assistive technology for cognitive rehabilitation: State of the art. *Neuropsychological rehabilitation* 14:5–39
30. Chan J, Nejat G. 2012. Social intelligence for a robot engaging people in cognitive training activities. *International Journal of Advanced Robotic Systems* 9:113
31. Tsiakas K, Abujelala M, Makedon F. 2018. Task engagement as personalization feedback for socially-assistive robots and cognitive training. *Technologies* 6:49
32. Woodworth B, Ferrari F, Zosa TE, Riek LD. 2018. *Preference learning in assistive robotics: Observational repeated inverse reinforcement learning*. In *Machine Learning for Healthcare Conference*, pp. 420–439
33. Rossi S, Ferland F, Tapus A. 2017. User profiling and behavioral adaptation for hri: A survey. *Pattern Recognition Letters* 99:3–12
34. 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. 359–366. IEEE
35. Tapus A, Ţăpuş C, Matarić MJ. 2008. User—robot personality matching and assistive robot behavior adaptation for post-stroke rehabilitation therapy. *Intelligent Service Robotics* 1:169
36. Organization WH, et al. 2017. Global action plan on the public health response to dementia 2017–2025

37. Teng E, Tassniyom K, Lu PH. 2012. Reduced quality-of-life ratings in mild cognitive impairment: analyses of subject and informant responses. *The American Journal of Geriatric Psychiatry* 20:1016–1025
38. Petersen RC. 2004. Mild cognitive impairment as a diagnostic entity. *Journal of internal medicine* 256:183–194
39. Bruscoli M, Lovestone S. 2004. Is mci really just early dementia? a systematic review of conversion studies. *International Psychogeriatrics* 16:129–140
40. Arnáiz E, Almkvist O. 2003. Neuropsychological features of mild cognitive impairment and preclinical alzheimer's disease. *Acta Neurologica Scandinavica* 107:34–41
41. 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
42. 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 research & therapy* 5:35
43. Garand L, Amanda Dew M, Eazor LR, DeKosky ST, Reynolds III CF. 2005. Caregiving burden and psychiatric morbidity in spouses of persons with mild cognitive impairment. *International journal of geriatric psychiatry* 20:512–522
44. 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*, pp. 1–15
45. Poulos CJ, Bayer A, Beaupre L, Clare L, Poulos RG, et al. 2017. A comprehensive approach to reablement in dementia. *Alzheimer's & Dementia: Translational Research & Clinical Interventions* 3:450–458
46. 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* 15:318–331
47. 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. 105–112
48. 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? *Universal Access in the Information Society* 4:105–120
49. Scassellati B. 2007. How social robots will help us to diagnose, treat, and understand autism. In *Robotics research*. Springer
50. 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. *Journal of autism and developmental disorders* 43:1038–1049
51. 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. *International journal of social robotics* 5:593–618
52. 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. *The Journal of Medical Investigation* 66:50–53
53. 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. *Schizophrenia research* 176:506–513
54. Giannopulu I. 2013. Multimodal cognitive nonverbal and verbal interactions: the neurorehabilitation of autistic children via mobile toy robots. *IARIA International Journal of Advances in Life Sciences* 5
55. 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. 3856–3859. IEEE

56. Eriksson J, Mataric MJ, Winstein CJ. 2005. *Hands-off assistive robotics for post-stroke arm rehabilitation.* In *9th International Conference on Rehabilitation Robotics, 2005. ICORR 2005.*, pp. 21–24. IEEE
57. Deng E, Mutlu B, Mataric MJ, et al. 2019. Embodiment in socially interactive robots. *Foundations and Trends® in Robotics* 7:251–356
58. 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. 3230–3235. IEEE
59. 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 geriatrics* 14:7
60. Graf B, Hans M, Schraft RD. 2004. Care-o-bot ii—development of a next generation robotic home assistant. *Autonomous robots* 16:193–205
61. 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. 2430–2437. IEEE
62. 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? *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 59:M83–M85
63. Yamazaki R, Kochi M, Zhu W, Kase H. 2018. A pilot study of robot reminiscence in dementia care. *International Journal of Medical, Health, Biomedical, Bioengineering and Pharmaceutical Engineering* 12:253–257
64. Tao C, Han R, Huang J, Wang X, Ma L. 2015. Development and experiment study of an intelligent walking-aid robot. *International Journal of Modelling, Identification and Control* 24:216–223
65. Libin A, Cohen-Mansfield J. 2004. Therapeutic robocat for nursing home residents with dementia: preliminary inquiry. *American Journal of Alzheimer's Disease & Other Dementias®* 19:111–116
66. 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. 101–102. IEEE
67. 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 (ICORR)*, pp. 1–6. IEEE
68. Woods B, O'Philbin L, Farrell EM, Spector AE, Orrell M. 2018. Reminiscence therapy for dementia. *Cochrane database of systematic reviews*
69. Wobbrock JO, Kane SK, Gajos KZ, Harada S, Froehlich J. 2011. Ability-based design: Concept, principles and examples. *ACM Transactions on Accessible Computing* 3:1–27
70. Morris MR. 2020. Ai and accessibility. *Communications of the ACM* 63:35–37
71. Mills M, Whittaker M. 2019. Disability, bias, and ai
72. 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. *Clinical interventions in aging* 13:863
73. 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* 278:1363–1371
74. 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. In *Mayo Clinic Proceedings*, vol. 89, pp. 1336–1349. Elsevier

75. 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. IEEE
76. Martins GS, Al Tair H, Santos L, Dias J. 2019. αpomdp: Pomdp-based user-adaptive decision-making for social robots. *Pattern Recognition Letters* 118:94–103
77. Charlton JI. 2000. *Nothing about us without us: Disability oppression and empowerment*. Univ of California Press
78. Lee HR, Riek LD. 2018. Reframing assistive robots to promote successful aging. *ACM Transactions on Human-Robot Interaction (THRI)* 7:1–23
79. Van Wynsberghe A. 2013. Designing robots for care: Care centered value-sensitive design. *Science and engineering ethics* 19:407–433
80. Castellano G, Aylett R, Dautenhahn K, Paiva A, McOwan PW, Ho S. 2008. *Long-term affect sensitive and socially interactive companions*. In *Proceedings of the 4th International Workshop on Human-Computer Conversation*
81. 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 Psychology* 34:1209
82. Fong T, Nourbakhsh I, Dautenhahn K. 2003. A survey of socially interactive robots. *Robotics and autonomous systems* 42:143–166
83. Szafir D, Mutlu B. 2012. *Pay attention! Designing adaptive agents that monitor and improve user engagement*. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 11–20
84. Yan H, Ang MH, Poo AN. 2014. A survey on perception methods for human–robot interaction in social robots. *International Journal of Social Robotics* 6:85–119
85. Pantic M, Pentland A, Nijholt A, Huang TS. 2007. Human computing and machine understanding of human behavior: A survey. In *Artifical Intelligence for Human Computing*. Springer
86. 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*, pp. 625–635. Springer
87. Meng Qm, Wu Wg. 2008. Artificial emotional model based on finite state machine. *Journal of Central South University of Technology* 15:694–699
88. 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. 303–307
89. Bohren J, Cousins S. 2010. The smach high-level executive [ros news]. *IEEE Robotics & Automation Magazine* 17:18–20
90. Liu C, Conn K, Sarkar N, Stone W. 2008. Online affect detection and robot behavior adaptation for intervention of children with autism. *IEEE transactions on robotics* 24:883–896
91. Sutton RS, Barto AG, et al. 1998. *Introduction to reinforcement learning*, vol. 135. MIT press Cambridge
92. 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. 265–266
93. Lockerd A, Breazeal C. 2004. *Tutelage and socially guided robot learning*. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, vol. 4, pp. 3475–3480. IEEE
94. Prommer T, Holzapfel H, Waibel A. 2006. *Rapid simulation-driven reinforcement learning of multimodal dialog strategies in human-robot interaction*. In *Ninth International Conference*

on Spoken Language Processing
95. Krsmanovic F, Spencer C, Jurafsky D, Ng AY. 2006. *Have we met? MDP based speaker ID for robot dialogue*. In *Ninth International Conference on Spoken Language Processing*
96. Malfaz M, Castro-González Á, Barber R, Salichs MA. 2011. A biologically inspired architecture for an autonomous and social robot. *IEEE Transactions on Autonomous Mental Development* 3:232–246
97. Ferreira E, Lefevre F. 2015. Reinforcement-learning based dialogue system for human–robot interactions with socially-inspired rewards. *Computer Speech & Language* 34:256–274
98. Karami AB, Sehaba K, Encelle B. 2016. Adaptive artificial companions learning from users' feedback. *Adaptive Behavior* 24:69–86
99. 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. *Computers & Education* 150:103836
100. 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 AAAI Conference on Artificial Intelligence*, vol. 33, pp. 687–694
101. Gao AY, Barendregt W, Castellano G. 2017. *Personalised human-robot co-adaptation in instructional settings using reinforcement learning*. In *IVA Workshop on Persuasive Embodied Agents for Behavior Change: PEACH 2017, August 27, Stockholm, Sweden*
102. Biancardi B, Mancini M, Lerner P, Pelachaud C. 2019. Managing an agent's self-presentational strategies during an interaction. *Frontiers in Robotics and AI* 6:93
103. 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 *Thirtieth AAAI Conference on Artificial Intelligence*
104. Wang Z, Singh MK, Zhang C, Riek LD, Chaudhuri K. 2020. *Stochastic Multi-Player Bandit Learning from Player-Dependent Feedback*. In *ICML Workshop on Real World Experiment Design and Active Learning*
105. Auer P, Cesa-Bianchi N, Freund Y, Schapire RE. 2002. The nonstochastic multiarmed bandit problem. *SIAM journal on computing* 32:48–77
106. 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*, pp. 1057–1063
107. Mitsunaga N, Smith C, Kanda T, Ishiguro H, Hagita N. 2006. Robot behavior adaptation for human-robot interaction based on policy gradient reinforcement learning. *Journal of the Robotics Society of Japan* 24:820–829
108. Smallwood RD, Sondik EJ. 1973. The optimal control of partially observable markov processes over a finite horizon. *Operations research* 21:1071–1088
109. 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. 544–549. IEEE
110. Hoey J, Von Bertoldi A, Poupart P, Mihailidis A. 2007. *Assisting persons with dementia during handwashing using a partially observable Markov decision process*. In *International Conference on Computer Vision Systems: Proceedings (2007)*
111. 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. 332–340. IEEE
112. Belpaeme T, Baxter P, Read R, Wood R, Cuayáhuitl H, et al. 2013. Multimodal child-robot interaction: Building social bonds. *Journal of Human-Robot Interaction* 1:33–53
113. Dietterich TG. 2000. Hierarchical reinforcement learning with the maxq value function decomposition. *Journal of artificial intelligence research* 13:227–303

114. Chandramohan S, Geist M, Lefevre F, Pietquin O. 2011. *User simulation in dialogue systems using inverse reinforcement learning*
115. Boularias A, Chinaei HR, Chaib-draa B. 2010. *Learning the reward model of dialogue pomdps from data*. In *NIPS Workshop on Machine Learning for Assistive Techniques*. Citeseer
116. Sugiyama H, Meguro T, Minami Y. 2012. *Preference-learning based inverse reinforcement learning for dialog control*. In *Thirteenth Annual Conference of the International Speech Communication Association*
117. Senft E, Baxter P, Kennedy J, Belpaeme T. 2015. *Sparc: Supervised progressively autonomous robot competencies*. In *International Conference on Social Robotics*, pp. 603–612. Springer
118. Dermouche S, Pelachaud C. 2019. *Generative Model of Agent's Behaviors in Human-Agent Interaction*. In *2019 International Conference on Multimodal Interaction*, pp. 375–384
119. Janecek A, Gansterer W, Demel M, Ecker G. 2008. *On the relationship between feature selection and classification accuracy*. In *New challenges for feature selection in data mining and knowledge discovery*, pp. 90–105
120. Gardner MW, Dorling S. 1998. Artificial neural networks (the multilayer perceptron)—a review of applications in the atmospheric sciences. *Atmospheric environment* 32:2627–2636
121. Hochreiter S, Schmidhuber J. 1997. Long short-term memory. *Neural computation* 9:1735–1780
122. 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)*, vol. 3, pp. 1–7. IEEE
123. Nigam A, Riek LD. 2015. *Social context perception for mobile robots*. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp. 3621–3627. IEEE
124. 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. *Applied Bionics and Biomechanics* 6:369–397
125. Sung JY, Guo L, Grinter RE, Christensen HI. 2007. *"My Roomba is Rambo": intimate home appliances*. In *International conference on ubiquitous computing*, pp. 145–162. Springer
126. 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. 258–265
127. Carpenter J. 2016. *Culture and human-robot interaction in militarized spaces: A war story*. Routledge
128. Riek L, Howard D. 2014. A code of ethics for the human-robot interaction profession. *Proceedings of We Robot*
129. 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. *Proceedings of We Robot*
130. Wang Z, Zhang C, Singh MK, Riek L, Chaudhuri K. 2021. *Multitask Bandit Learning Through Heterogeneous Feedback Aggregation*. In *International Conference on Artificial Intelligence and Statistics*, pp. 1531–1539. PMLR
131. Auer P, Cesa-Bianchi N, Fischer P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47:235–256
132. Kato Y, Kanda T, Ishiguro H. 2015. *May I help you?-Design of human-like polite approaching behavior*. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 35–42. IEEE
133. 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. 518–521. IEEE
134. Spaulding S, Breazeal C. 2017. *Learning behavior policies for interactive educational play*. In *Models, Algorithms, and HRI Workshop*
135. Chu MT, Khosla R, Khaksar SMS, Nguyen K. 2017. Service innovation through social robot engagement to improve dementia care quality. *Assistive Technology* 29:8–18