



Challenges in Designing a Fully Autonomous Socially Assistive Robot for People with Parkinson's Disease

JASON R. WILSON, Department of Computer Science, Northwestern University

LINDA TICKLE-DEGNEN, Department of Occupational Therapy, Tufts University

MATTHIAS SCHEUTZ, Human-Robot Interaction Lab, Tufts University

Assistive robots are becoming an increasingly important application platform for research in robotics, AI, and HRI, as there is a pressing need to develop systems that support the elderly and people with disabilities, with a clear path to market. Yet, what remains unclear is whether current autonomous systems are already up to the task or whether additional HRI work is needed to make these systems acceptable and useful.

In this article, we report our efforts of developing and evaluating an architecture for a fully autonomous robot designed to assist older adults with Parkinson's disease (PD) in sorting their medications. The main goal for the robot is to aid users in a manner that maintains the autonomy of the user by providing cognitive and social support with varying levels of assistance. We first evaluated the robot with subjects drawn from a pool of university students, which is common practice in experimental work in psychology and HRI. As the results were very positive, we followed up with an evaluation using people with Parkinson's disease, who surprisingly had mostly negative outcomes. We thus report our analysis of the differences in the evaluations and discuss the challenges for HRI posed by the sources of the negative evaluations: (1) designing a robot to adapt to the many routines the participants use at home, (2) unique needs of participants with PD not present in student participants, and (3) the role of familiar technologies in designing and evaluating a new technology. While it is unlikely, given the current state of technology, that fully autonomous assistive robots for older adults will be available in the near term, we believe that our work exposes a critical need in HRI to involve the target population as early as possible in the design process.

CCS Concepts: • **Human-centered computing** → *Natural language interfaces*; • **Computing methodologies** → *Cognitive robotics*; *Robotic planning*; • **Social and professional topics** → *Assistive technologies*;

Additional Key Words and Phrases: Social robot, assistive technology, robot architecture, cognitive robot

ACM Reference format:

Jason R. Wilson, Linda Tickle-Degnen, and Matthias Scheutz. 2019. Challenges in Designing a Fully Autonomous Socially Assistive Robot for People with Parkinson's Disease. *ACM Trans. Hum.-Robot Interact.* 9, 3, Article 20 (May 2020), 31 pages.

<https://doi.org/10.1145/3379179>

This work was funded by NSF (Grant No. IIS-1316809). The authors declare that they have no conflict of interest.

Authors' addresses: Jason R. Wilson, Department of Computer Science, Northwestern University, 2233 Tech Drive, Evanston, IL 60626; email: jrw@northwestern.edu; L. Tickle-Degnen, Department of Occupational Therapy, Tufts University, 574 Boston Ave., Medford, MA 02155; email: linda.tickle_degnen@tufts.edu; M. Scheutz, Human-Robot Interaction Lab, Tufts University, 200 Boston Ave., Medford, MA 02155; email: matthias.scheutz@tufts.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-9522/2019/05-ART20

<https://doi.org/10.1145/3379179>

1 INTRODUCTION

As people grow older, they inevitably face many challenges in managing their own health and could thus benefit from getting assistive devices like medication management systems. Yet, most older adults prefer to be involved in health decisions [27] and want to take an active role in “designing their own lives” [52]. As we explore how assistive technologies, in particular, assistive robots, can contribute to health management activities, we must consider how technology can promote an active role of the persons being assisted, support the persons’ autonomy while being assisted, and enable people to care for themselves. Achieving medication adherence, for example, is a significant challenge for many older folks who must follow a daily regimen of medications, which is defined as the extent to which patients take medications as prescribed [34]. It requires successful “medication management”—the ability to develop, schedule, and implement a plan to take medications and know if medications have been taken [49]. Successfully adhering to a medication management schedule is often critical for managing and delaying the progression of chronic diseases, such as Parkinson’s disease.

In this work, we investigate how a socially assistive robot may help a person with Parkinson’s disease manage medications by assisting in the sorting of medications. Parkinson’s disease (PD) is a neurodegenerative condition, for which there is currently no cure, but symptoms can be dramatically reduced with antiparkinsonian medication. Optimizing adherence to oral medication regimens is critical to managing the symptoms of PD [13, 22], but many challenges impede medication adherence [45].

To improve medication adherence, we propose that a social robot may assist a person in sorting and organizing medication. Others have tried incorporating a variety of other assistive devices. Possible solutions are pill holders, alarm-based pill holders, and pill monitoring devices [38]. Pill holders are *passive solutions* that organize medications according to day and/or time. Augmented pill holders use sensors to track information like when compartments are opened or closed [16, 26]. In contrast, *active solutions* like alarms and pill-monitoring devices alert a person [11, 54] and dispense medications [5, 32, 42] when it is time to take a medication. Active solutions with robots may also provide medication reminders [10, 56] or may physically deliver the medication to the user [47]. While these solutions are active from the perspective of the technology, the user plays mostly a *passive role*. Instead, we would like the user to take a more active role as solutions that allow the user to take a passive role, such as pill boxes and automated dispensers, have been found to be less engaging, provide insufficient instructions and education, and are unable to provide affective and social support [51]. Alternatively, a social robot has the potential to be more engaging, provide instruction, and provide affective and social support [53], e.g., using non-verbal behaviors to encourage user engagement [46, 58]. Through user engagement and verbal interaction, a social robot may also educate the user, enabling the user to have greater knowledge about the medications and reasoning behind the prescriptions. Importantly, education and user knowledge have been shown to contribute to better medication adherence [9, 29].

However, a social robot that is fully autonomous (as will be necessary when it is deployed in one’s home) brings far more technological challenges than the simpler solutions mentioned above. Some have found that it is too challenging to reach the necessary robustness to have a fully automated system that uses natural modes of interaction like speech, and as a result touch screens and other (non-social) forms of interaction are required for interacting with the robot [36].

In this work, we present the architecture for a fully autonomous social robot designed to assist in medication sorting. We then first describe the evaluation performed with student subjects, which was used to aid in the development of the system, to provide evidence that the technology is working as expected, and that it meets the users’ expectations. Following the positive outcome,

we then report an evaluation using our target population, people with PD who are more vulnerable (due to age and condition) and more difficult to enroll (due to their disease and low numbers of available subjects). Despite the strongly positive evaluation scores with the students, the latter evaluation did not fare well, and most participants provided a mostly negative evaluation of the system. We discuss later in the article some of the reasons why older adults responded negatively to the interaction with the robot and highlight some of the challenges in deploying a fully autonomous robot to assist older adults, which stem from the fact that older adults have different needs than younger adults and have generally less familiarity with technology. Moreover, sorting medication is more relevant to their lives, and, as a result, adults with PD evaluate the robot from a different (and more valuable) perspective. While the exposed challenges are not insurmountable, they do provide barriers for deploying social robots to assist older adults in the near future, and we conclude that it is important for HRI research to start involving target populations (rather than university student subjects) as early as possible in the design process.

2 BACKGROUND ON ASSISTING IN MEDICATION MANAGEMENT

2.1 Challenges in Medication Management

Adherence to a medication regimen is critical to treatment outcome and quality of life for people with PD. However, they face many challenges to achieving medication adherence. In addition to general issues elderly patients may have (e.g., age-related physical decline, economic factors), people with PD often require complicated dosing or titration schedules and may have co-morbidities that require the coordination of therapies from multiple drug classes. Disease progression can also introduce cognitive impairments that can affect adherence, responses to antiparkinsonian agents can cause variable responses interfering with medication adherence [4].

One effective approach to achieving medication adherence is to use aids such as pill boxes and hour-by-hour organizational charts [4]. Managing medications with a pill box requires a person to sort medications into the box, which entails organizing pills according to the day of the week and time each pill is to be taken. Each pill is placed into a pill box or on a sorting grid, as shown in Figure 1. The grid contains columns for each day of the week and rows for different times of day, such as Morning, Noon, Evening, and Bedtime.

While using pill containers are suggested for improving medication adherence [4, 29], sorting medications can be a cognitively demanding task requiring verbal memory, cognitive flexibility, and executive functioning [50]. For people with PD, managing medications can be even more challenging, as they may require a more complicated medication regimen that can include antiparkinsonian medication, medications for co-morbidities, and a number of over-the-counter supplements and herbal preparations. Additionally, disease progression can introduce memory disorders and cognitive impairments that can affect adherence [4]. As a result of these numerous challenges, people with PD have shown performance deficits in managing medications and may require more assistance as compared to other older adults [12]. Though they may need more assistance, one needs to still be careful in how to provide assistance, to not be too controlling or over-bearing, as medication decisions is one activity in which older adults desire to maintain autonomy [15].

2.2 Assisting while Supporting Autonomy

To help develop and maintain the skills necessary for managing medications and other self-care activities necessary for daily living, a person with PD may be referred to an occupational therapist (OT). The OT works with the client to assess the person's needs, develop an intervention plan, and help the person achieve his or her goals. The intervention plan may include how to incorporate assistive technologies (including robots) to maintain or improve the client's capabilities.

SUNDAY	MONDAY	TUESDAY	WEDNESDAY	THURSDAY	FRIDAY	SATURDAY
Morning	813:0	Morning	Morning	826:0	Morning	Morning
Afternoon	Afternoon	Afternoon	Afternoon	Afternoon	Afternoon	Afternoon
809:0 810:0	Evening	824:0	Evening	Evening	Evening	Evening
Bedtime	Bedtime	Bedtime	Bedtime	Bedtime	Bedtime	Bedtime

Fig. 1. The medication grid has columns for each day and four rows for morning, noon, evening, and bedtime. The Vision system recognizes pills and identifies their locations within the grid.

In the context of medication management, this may include devices to provide reminders (e.g., [11, 54]) or dispense medications (e.g., [5, 32, 42]). There are fewer options available for organizing medications. One option is PillPack, which pre-sorts and packages medications into individual packets [35]. While this solution may be appropriate for some, pre-sorted packets removing the decision-making responsibility from the person taking the medications, and many older adults want to maintain autonomy in medication treatment decisions [15].

When an OT assists a person, one overarching goal is to support the autonomy of the person by respecting individuals' choices and empowering clients to exercise their right to make decisions based on the direct care of their own health [3]. *Autonomy* has been described as regulation by the self [57], and we extend this definition to describe autonomy as a continuum corresponding to the degree to which one is free to make choices and act in accordance with one's goals, needs, and wants. A person's sense of autonomy may be influenced by how much a person is able to choose a course of action [19], be in control [19, 25, 57], have the necessary resources and knowledge to take actions [19], and be the author of his or her actions [57]. Conversely, one can inhibit a person's autonomy by preventing the person from doing what he or she wants to do [33].

While there are many influences of autonomy for people of all ages (e.g., injury, imprisonment, hindering), older adults face additional challenges due to age and illness [33]. As adults age, they may become dependent on others to do or assist with daily activities, and as a result, they are at risk of being deprived of their autonomy [37]. However, dependence does not necessarily imply a complete loss of autonomy if the dependence does not strip them of their right and ability to freely act and make decisions. Instead, assistance needs to enable the individual, thus allowing older adults to be autonomously dependent while still needing to rely on others [19, 41]. By maintaining their autonomy, they can preserve their dignity [33, 59] and are more likely to exhibit higher morale and greater life satisfaction [19].

We envision a social robot giving assistance while maintaining the autonomy of the older adult being assisted, and we examine how an OT approaches assisting a client to gain insights in how a




Fig. 2. The nine levels of assistance as defined in the PASS manual describes increasing amounts of assistance.

robot can support the autonomy of an older adult. In working with a client, an OT needs to regulate how much assistance is provided to not control the person, dictate each of their actions. Providing constant assistance may ensure that the task the client is doing is done correctly and safely, but it could also divest the client of their dignity and autonomy. The client should be able to have the opportunity to choose each action he or she takes in a given task, and thus too much assistance often contributes to a loss of autonomy. However, providing too little assistance can also have a negative impact, as it deprives the client of the necessary resources to do the task, which impacts the client's sense of autonomy [19].

Instead, the OT needs to match the amount of assistance provided to the amount of need the person has. To guide an OT in providing graded assistance, one may rely on a hierarchy of levels of assistance, as defined in the Performance Assessment of Self-care Skills (PASS) manual [40], which describes nine levels of assistance (see Figure 2).

The first few levels of assistance rely on non-physical support and describe verbal assistance. In Section 3.4.2, we describe how our socially assistive robot uses the first four levels of assistance to provide cognitive and social support to assist a person in sorting medications while attempting to preserve the autonomy of the individual.

2.3 Task Analysis

In addition to understanding how assistance relates to autonomy, we also need a detailed description of the task with which a robot is to assist. To better understand medication sorting procedures, the steps involved, and the complexities of managing multiple prescriptions, we developed a thorough task analysis. We have previously reported this analysis elsewhere [60], and we summarize the important outcomes here.

We developed a script of a typical medication sorting task based on a standardized protocol defined in [40]. Based upon this script, we recorded a video that simulated an execution of a medication sorting activity. In the video, one actor simulated a person with PD and placed two medications onto a sorting grid. The actor read the instruction of the first medication and then sorted the medication by placing one pill at a time. The same procedure was followed for the second medication. Meanwhile, another actor, who simulated the role of a caregiver, guided the first person by providing verbal cues and acknowledgments to progress the task. The script and video were then used to develop a task analysis [39], which captured the objects used, environmental demands, social demands, required actions, prerequisite capacities, and safety concerns. The video and task




Fig. 3. The Vision, Speech, and Sensory Integration components compose the “Sense” portion of the architecture. The “Think” components are the Goal Manager and the Medication Management and Assistance components. The Nao and Soundboard components are the “Act” components.

analysis were then reviewed by a focus group of four occupational therapists with experience working with people with PD.

The task analysis was found to be comprehensive and complete, but the video lacked some of the complexities of a realistic setup. For example, in the video, the person simulating the role of the person with PD poured the pills onto the table, but it was suggested they would instead pour it into a bowl to prevent the pills from rolling off the table. The focus group also indicated that the sorting grid should be positioned directly in front of the person sorting the medications.

Based on these results, we designed the robot to assist in a medication sorting task as described in this task analysis (see Section 3.4.2). An important design goal in designing a robot to assist with activities such as medication management is to ensure the robot can contribute to preserving the dignity and maintaining the autonomy of the person being assisted [2, 59]. In the work described here, we designed a system that provides cognitive and social support to a person sorting medications. Instead of providing direct physical support to the person, the robot provides assistance that is intended to enable the person to make the correct actions to successfully sort the medications. The assistance the robot gives comes mostly in the form of spoken assistance, but the robot can also change its gaze to look at the person and make simple gestures like pointing. The robot is able to follow the course of the task by visually monitoring the task and listening to the speech of the person.

3 ROBOT ARCHITECTURE

This section describes the cognitive robotic architecture used to control the fully autonomous robot that assists in sorting medications. The components of the architecture provide the mechanisms by which the robot can see and hear its environment, decide how to react to stimulus, and then act within its environment.

3.1 High-level Architecture

The seven components of the system were implemented in ADE [43], the implementation middleware of the Distributed, Integrated, Affect, Reflection, Cognition (DIARC) architecture [44]. Each of the components can be classified to correspond to one of the steps in the classical “sense-think-act” cycle (see Figure 3). The components for interacting with the environment (i.e., sense and act

components) are described first, and then we describe the core functionality of the autonomous robot in the section on the think components.

3.2 Sense

The sensory components interpret video and audio information to produce a representation of the perceived state of a task and the person performing that task. The Vision component processes the video data, the Speech component processes the audio data, and the Sensory Integration component combines the results of each of these components into a single event to be processed by the "Think" components.

3.2.1 Vision. The Vision component provides processing of the video feeds from the robot cameras to determine information about objects in the environment and about the person with whom the robot is interacting. In perceiving its environment, the robot can use the Vision system to recognize and locate objects, identify properties of the object (e.g., color), and determine if another object (such as a hand) is occluding the object. In reference to a person, it can infer the direction of the person's gaze.

To observe any progress in the sorting of medications, the robot needs to be able to monitor the medication grid and recognize any changes. To make visual observations of the grid, the robot has a Vision component that uses the cameras on the robot to capture the state of the medication sorting grid (shown in Figure 1). The Vision system reports how many pills of each type are in each cell in the grid, where the type of pill is determined by the pill color. A change in the number of pills in any of the cells of the grid is later used by the Medication Management and Assistance component to infer whether an action might have been taken by the person (see Section 3.4.2).

In addition to observing the sorting grid, the Vision system is also responsible for observing the person that is doing the medication sorting. In particular, it recognizes when the gaze of the person shifts from the grid to the robot or from the robot to the grid.

For the robot to simultaneously observe both the sorting grid and the direction of the person's gaze, the robot uses two cameras (see Figure 4). The bottom camera tracks the sorting grid while the top camera tracks the person's gaze.

3.2.2 Speech. The Speech component provides automatic speech recognition so that the robot can determine if the person needs help with the task. A person may ask about scheduled events (e.g., "When is physical therapy?"), or the person may conclude that the task is done and say, "I'm done." The robot needs to be able to interpret these utterances and respond accordingly. For example, if the person says "I'm done," but there is a missing pill, then the robot needs to be able to indicate that there are more pills to be sorted.

One goal for the Speech component is to allow for free, unstructured speech to create a more natural interaction with the robot. In other words, predefined phrases that a person would need to know are avoided. This is especially a concern for older adults, as we felt it was important to not require any sort of training to be assisted by the robot and that the medication sorting would ideally be done similarly to how the person would do it without the robot's assistance.

Instead of fully recognizing specific phrases, the Speech component only needs to identify a few keywords. This means that not all of the speech needs to be fully recognized, and the utterance does not need to be parsed. These simplifications provide greater flexibility in terms of what can be recognized.

By recognizing keywords, the system would be able to determine whether the person is asking a question, references a day of the week, or makes a comment about being done. Examples include the following (with keywords emphasized): "*When* do I have physical therapy?", "Is *Monday* ok?",




Fig. 4. The Nao robot has two cameras. The top is used to monitor the gaze of the person. The bottom is used to track the pills on the sorting grid and tray.

and “I’m *done*.” These keywords are later used by the Medication Management and Assistance component (see 3.4.2) to determine the appropriate type of action.

The speech recognition is accomplished using Sphinx4 [23], which is configured to recognize utterances from a user defined grammar in JSGF format.¹ The intention of the grammar is to loosely define the structure of utterances, focusing on the relations between a small set of keywords. The complete set of rules in the grammar is the following:

```
public <utterance> = <wQuestion> |[ok] [i'm | we] done [yet] | <day> | <help> |<extra>;
<help> = [i] [need] help [me] [with];
<extra> = yes | no | ok | hello | hi;
<wQuestion>= what | when | where | which| why;
<day>= day | sunday | monday | tuesday | wednesday | thursday | friday | saturday;
```

To handle utterances not defined in our grammar, we add context independent phone loops to the decoder graph. This modification allows the recognizer to transition to the “unknown word” state at any point during decoding, including at the phone level in the middle of words [17]. This enables the recognizer to identify when a word or phrase is not in its grammar. As a result, the output of the recognizer is still limited to the set of utterances defined in its grammar, but in cases where the input does not match the grammar, it is much less likely to produce a false positive.

3.2.3 Sensory Integration. Sensory Integration ensures that downstream processing always has a complete and accurate representation of the perceptual information the robot receives. The Vision component is constantly processing the video feed and always has a representation of the

¹<https://www.w3.org/TR/jsgf/>.

current state of the grid and the gaze. The Speech component processing is intermittent, only recognizing utterances when the person says something. The responsibility of the Sensory Integration component then is to include the state from the Vision system in every event that the Sensory Integration component sends, and recognized speech utterances are included in events only when they occur.

To accomplish this, updates from Vision and Speech are handled slightly differently. When a new event from Vision is received, this is immediately passed along to the Goal Manager. When a new recognized utterance is received from Speech, the Sensory Integration component will query the Vision component for the most recent visual state, and then package the Vision and Speech events together and send them along.

To communicate with the Goal Manager, the Sensory Integration component submits a goal to the Goal Manager (see 3.4.1 for details of how the Goal Manager handles goals). The goal given to the Goal Manager is the predicate `handleMedicationEvent(e)`. The argument to this predicate is the event that encapsulates the details of the information gathered by the Vision and Speech components.

3.3 Act

In the DIARC architecture, behaviors for the robot to perform are represented as *actions*, where there are two types of actions: *primitive actions* and *action scripts*. An action script is a composition of actions, specifying some sequence of actions to take and control logic managing the flow of actions. A primitive action is an atomic action that may be executed, and it maps to a procedure call provided by an “Act” component controlling the robot.

We define ten types of actions that the robot may select to assist a person in medication sorting:

- `sayText`
- `blink`
- `closeHand`
- `goToPosture`
- `pointHeadTo`
- `pointTo`
- `noOp`
- `lookAndSay`
- `pointToError`
- `notifyDone`

The first six are primitive actions used to control the Nao robot and are described in Section 3.3.1. The seventh action (`noOp`) represents a lack of action and is used when the robot should not respond to what it has observed.

The last three actions are action scripts, each of which decomposes into additional actions. The `lookAndSay` action, which is composed of actions `pointHeadTo` and `sayText`, is for when we want the robot to look at the person while speaking (as opposed to looking down at the medication grid). The `pointToError` action makes the robot point to a location on the medication grid while describing the error it observes. This action script is composed of actions for `sayText` and `pointTo`. The `notifyDone` action notifies the person that the medication (or all medications) have been successfully sorted by executing a `sayText` action. The `notifyDone` action has a precondition defined so that the action is executed only when medications are done being sorted.

In the rest of this section, we describe the two “Act” components. The primary component is the Nao component, which provides *primitive actions* that encapsulate the low-level mechanisms to control a Nao robot. The component may be used to control a physical Nao robot, as is customarily




Fig. 5. The physical and virtual Nao setups were similar, with the size and dimensions of the robot nearly matching.

done, and it can also control a virtual Nao robot (used in our first evaluation, see Section 4.1). The virtual Nao runs as a simulation in the *Choregraphe* tool (see Figure 5 for images of the physical and virtual robots). Since *Choregraphe* does not support audio production of speech, we introduced the Soundboard component to handle the speech production of the virtual Nao.

3.3.1 Nao. The Nao component provides the interface between the DIARC architecture and the control software for the Nao robot. The interface hides the low-level mechanics of joint positions and movement velocities, instead provides an interface that is more abstract, such as pointing at an object. It also provides direct access to the robot's text-to-speech facility.

For the purposes of medication sorting, the following actions provided by the Nao component are used:

- `blink`: make the robot's eyes appear to blink
- `closeHand`: close one of the robot's hands
- `goToPosture`: move to a preset posture
- `pointHeadTo`: move head to be directed at a point
- `pointTo`: move arm to be directed at a point
- `sayText`: synthesize speech output

Since the interface to the physical Nao and the virtual Nao (running in *Choregraphe*) are the same, this component is used to control both embodiments of the robot. However, the virtual Nao is not capable of speech production. In this case, the `sayText` interface to this component is removed so that the Soundboard component may handle speech production.

3.3.2 Soundboard. The Soundboard component provides speech capabilities to the virtual Nao. Similar to the Nao component, it uses the same `sayText` interface to produce the speech output. Unlike the Nao component, which uses the underlying Nao API to synthesize speech, the Soundboard component does not synthesize speech. Instead, it plays an audio file that has been mapped to a particular utterance. To ensure that the speech heard by the user is the same for both the physical and virtual robots, the physical robot was used to record each possible utterance to an audio file. A built-in feature of the Nao API allowed us to redirect the robot's audio output to an audio file, which gave us a clear reproduction of the robot's voice.

3.4 Think

The main “Think” component in the DIARC architecture is the Goal Manager, which allows the robot to autonomously decide how to respond to its observations. As part of the robot's thinking process, the Goal Manager may rely on other components to assist in the decision-making process, and in the case of medication sorting, the domain-general Goal Manager is assisted by the domain-specific Medication Management and Assistance component. Each of these components are described in this section.

3.4.1 Goal Manager. The Goal Manager component is responsible for managing the goals of the robot and the actions taken to satisfy these goals. It provides much of the core functionality that enables the robot to act without the intervention of any remote operators.

Processing of goals happens in three steps:

- (1) A goal is submitted to the Goal Manager.
- (2) An action that will accomplish this goal is selected.
- (3) The action is executed.

A *goal* is defined with a predicate that specifies the end state that the robot needs to achieve. For example, the robot may be given the goal `medsSorted()`, indicating that the robot will have achieved this goal once the medications have been sorted. To achieve this goal, the robot may select and execute an action. There are two types of actions: *primitive actions* and *action scripts*. An action script is a composition of actions, specifying some sequence of actions to take and control logic managing the flow of actions. A primitive action is an atomic action that may be executed.

There are two goals relevant to assisting in sorting medications. One is given to the robot at startup and the other is given to the robot in response to the robot's observations. The first goal is `medsSorted()` and informs the robot that it needs to assist with the medication sorting. The second goal is `handleMedicationEvent(e)` and instructs the robot to process the new stimulus data to assess the progress in the medication sorting and respond accordingly. The first goal remains active until the medication sorting is complete, and the robot pursues the second goal in parallel with the first.

In response to the `medsSorted()` goal, the Goal Manager selects the top-level Action Script 1, which specifies the overall flow of the robot's actions while it assists a person that is sorting medications. Each of the steps in this action script is another action script. The first step is the action `instructMedSorting`, which defines steps for the robot to greet the person (e.g., saying “Welcome”), inform the person that the robot will assist in medication sorting, and orient the person to the task by pointing out the sorting grid and the pill bottles. The second and last actions start and shutdown the processing in the Vision component.

The third step provides the core actions for assisting with medications. The steps of this action script are shown in Action Script 2. The action script starts by having the robot provide some instructions to the person sorting the medications. Then the robot is to look down, detect the grid, and provide one more instruction. The remainder of the script has it looping to check if all of the pills have been correctly sorted. Inside that loop is checking for new events and then processing the new events. The `hasNewEvent` is true when the Sensory Integration component has received a new event from the Speech or Vision components.

During the execution of the `processNextEvent` action, a new goal `handleMedicationEvent(e)` is submitted, which instructs the robot to handle the most recent sensory data. The action(s) the robot takes in response to this goal provide cognitive and social assistance to the person that is sorting the medications. To select the actions for this goal, the Goal Manager relies on the Medication Management and Assistance component.

Action Script 1: medSortingTask()

*Top-level script for robot assisting in medication sorting***PostConditions**

medsSorted()

EndPost**Steps**

instructMedSorting

readyVision

assistMedSorting

cleanupVision

EndSteps

Action Script 2: assistMedSorting()

*Top-level script for robot assisting in medication sorting***Steps**

sayText "Pick one of the containers of pills."

sleep 1500

sayText "Read the instructions on label out loud and
then pour some on the tray."

lookAtGrid

detectGrid

sayText "When you are ready, start placing the pills
in the sorting grid."**while** not(isDone) **do** **while** not(hasNewEvent) **do**

sleep 500

end while

processNextEvent

end while**EndSteps**

3.4.2 Medication Management and Assistance. When assisting in medication sorting, the Goal Manager uses the Medication Management and Assistance component to select the actions to achieve goals related to medication management. In calling out to the Medication Management and Assistance component, the Goal Manager uses a general mechanism whereby it may call out to a domain reasoner to select an appropriate action. The Medication Management and Assistance component reasons about how much need for assistance the person has, what type of assistance is needed, and selects an appropriate action. For example, if a pill has been placed on a Tuesday but it should be on a Wednesday, and if it is determined that the person needs clear information about the misplaced pill, then the component may select an action that has the robot point at the misplaced pill and say "The blue pill on Tuesday is not needed." However, if the person does not need that much information, then it may suffice for the robot to not point and simply say "Make sure everything is on the right day," which would allow the person to find the misplaced pill on his or her own.

Of the ten actions defined for controlling the Nao robot (see Section 3.3), five of these actions may be selected by the robot when assisting a person in medication sorting: sayText, noOp,

lookAndSay, pointToError, and notifyDone. The other five actions are only available indirectly via an action script. The rest of this section only references the five actions that may be directly selected when deciding how best to provide assistance.

In addition to the five primitive actions, the robot may select one of three action scripts: lookAndSay, pointToError, and notifyDone. The lookAndSay action script is for when we want the robot to look at the person while speaking (as opposed to looking down at the medication grid). This is typically selected when the robot has detected that the person's gaze is directed away from the grid and up at the robot. This is done so that the robot responds to the person while looking at the person when the person is looking at the robot. The pointToError action script is used for the level 4 assistance (gesture) and has the robot point to the location on the grid where a misplaced pill has been detected. The notifyDone action script notifies the person that the medication (or all medications) have been successfully sorted, but its more important function is to advance the action script of the overall task so that the robot can either provide new instruction to the person or indicate that the task is complete.

To select which action is most appropriate, the component takes the following four steps:

- (1) Preprocessing the event.
- (2) Categorize the state of the task.
- (3) Select an assistance for that category.
- (4) Update the estimated need the person has for assistance.

Step 1. Preprocessing. To supplement the sensory data in the event (i.e., state of the medication grid, user utterance, direction of gaze), the preprocessing step adds information regarding

- misplaced pills,
- task completion,
- presence of a user action,
- correctness of a user action,
- user hesitation or delays, and
- quantity of recent errors.

The preprocessing step analyzes the sensory data to compile the additional information so that it may be readily available for the following steps.

To identify any misplaced pills or to determine whether the task is complete, the state of the medication grid (with the quantity of each pill type in each location of the grid), is compared with the desired goal state of the medication grid. If the current grid state perfectly matches the goal state, then the task is complete. If any cell of the grid in the current state has more pills of a particular type than is specified in the goal state, then the current state is inconsistent with the goal, and the location of the misplaced pill is recorded.

To identify whether the person has taken some action, a comparison between the states of the current grid and previous grid reveal whether any pill has been added, moved, or removed. Changes in the state of the grid are assumed to be caused by the person doing the task, and thus a user action is inferred. If there is no change in the grid, then the event the robot needs to respond to is the result of either a recognized speech utterance or a change in the direction of the person's gaze (e.g., the person was looking at the grid but is not looking at the robot).

If the component has inferred that the person has made an action, then it needs to determine if the action was a correct one or not. An action is considered correct if (1) a pill has been *added* to a position in the grid and the number of pills now in that position is less than or equal to how many the goal state indicates should be there, and/or (2) a pill has been *removed* from a position




Fig. 6. The decision tree for classifying each state has yes answers on the left and no answers on the right of each decision node. The boxes at the leaves represent the nine categories.

in the grid and the number of pills previously in that position is greater than how many the goal state indicates should be there.

To recognize a possible hesitation on the part of the user, the system examines how long it has been since the last change in the state of the grid. A user hesitation is inferred if there has been no change for four seconds. During this time, the user may be speaking or looking at the robot, looking at the pill bottles, but is not adding, removing, or moving any pills on the medication grid.

Last, the system calculates how many “recent” errors have been made. To determine this, the component scans over the recent events to find how many of them represent some form of error. An error is either an incorrect action or a hesitation. An event is considered “recent” if it was one of the five most recent events. Thus, the number of recent errors is the number of events out of the past five that represent an error.

Step 2. State Categorization. The state of the task needs to be categorized to determine what type of assistance needs to be provided. For example, if the preprocessing determines that a pill has been added and that the pill placement was incorrect, then the robot needs to give some type of assistance particular to fixing a mistake. To classify the state of the task, we constructed a large binary decision tree. The upper portion of the tree classifies the state into one of nine categories (see Figure 6).

The first five categories (reading the figure left to right) are related to there being no misplaced pills on the medication grid. The *Done* categories indicate that either all the pills for a single medication has been completely sorted or all the medications have been completed and the task is done. The *Good* category means a pill has been added or removed from the grid and that all pills on the sorting grid have been correctly placed. *Incomplete* is when there has been no change in the state of the grid, there are no misplaced pills, and the task is not complete. Since there has been no change in the state of the grid, the robot is responding to either something the person has said or a change in the direction of the person’s gaze. The *Not started* category is for the beginning of the task and handles the case when the person is not sure whether or how to begin.

The next four categories indicate that there is at least one mistake represented in the current state of the medication grid. The *Wrong day* category is for when a pill is placed on a day in which the pill should not to be taken. The *Wrong time* category is for when a pill is placed on the correct

day but at the wrong time. A special case of the *Wrong time* category is when a pill is not only at the wrong time but all of the correct times are fully allocated. For example, if a pill is to be taken on Sunday with breakfast, and a person has one pill in the morning cell and another one in the evening, then this was a *Wrong time and too many* event.

Once the top level category has been determined, then the decisions in the tree further categorize the state. The types of decisions include determining which medication had been misplaced, whether the person is looking at the robot or not, and recognizing whether the person said something and whether the content of the spoken utterance included a particular keyword. For example, for an event categorized as *Incomplete*, the decisions that follow included checking whether a speech utterance was detected, if the utterance includes the word “done,” and which medication is incomplete. Then the selected assistance is an action that informs the person about which medication is incomplete.

Step 3. Assistance Selection. Once the state has been categorized, then the Medication Management and Assistance component determines the appropriate level of assistance. Given the level of assistance, it selects an action that addresses the current state of the task.

Instrumental to selecting the appropriate assistance is a taxonomy of levels of assistance. The *Performance Assessment of Self-care Skills* (PASS) manual, which is used in occupational therapy practice, defines nine levels of assistance [40]. In focusing on cognitive and social support rather than physical support, we limited the robot to only being able to provide the first four levels of assistance. Starting at level 5, the assistance may require physical contact with the person being assisted. Limiting the robot to the first four levels emphasizes the cognitive and social support that the robot is designed to provide.

For the robot to provide assistance of a particular level, how the actions of the robot relate to assistance needs to be defined. We describe here each of the levels of assistance that the robot is able to provide and examples of actions the robot may take to provide that assistance.

Level 0. A level 0 assistance is used when the robot provides information or instruction about the task. The PASS does not define a level 0, but it was necessary to categorize the actions the robot takes that is not direct assistance to the task but are simply informing the person about the task. There are two examples of this: (1) the robot may indicate to the person that he or she may begin sorting, or (2) the robot notifies the person that one medication has been successfully sorted.

The robot may indicate that the person can start sorting the pills if no pills have been placed yet and there has been a long delay since the robot completed its initial instructions. In this case, the robot is to perform the action `sayText(“You may begin sorting the pills now”)`.

When the last pill of a medication has been correctly placed, the robot notifies the person that he or she has successfully completed sorting that medication. Similarly, once all the pills have been successfully sorted, the robot indicates it. Both of these actions are accomplished via the `notifyDone` action script, which uses the `lookAndSay` action to look at the person and either say that the task is complete or say that the green or blue pills are done.

Level 1. A level 1 assistance is intended to be *verbal supportive*, providing encouragement for the person to continue or to complete the task [40]. This assistance is selected when the person does not currently need much assistance, perhaps because he or she has made few, if any, mistakes. Often, a level 1 assistance is simply the robot saying “Ok” and informs the person that everything is going well. Having the robot make this short utterance also lets the person know that the robot is still functioning and still following along. Another case where a level 1 assistance may be necessary is when not all the pills have been sorted yet, but the person may believe that he or she is done. A level 1 assistance in this case results in the action `sayText(“Check the grid carefully.”)`.

The utterance that is then spoken by the robot does not provide any indication as to whether there is anything right or wrong and only encourages the person to check things over. However, if the person does not see what is missing or does not realize that more pills need to be sorted, then more assistance may be needed.

Level 2. The PASS defines a level 2 assistance as *verbal non-directive* [40] and is used to provide some verbal indication that an action on the part of the person is necessary without directly saying what action needs to be taken. For example, in the case where the task is incomplete, as described in the previous level, the level 2 assistance is the action `sayText('I think there's more to be done.')`. When the robot says this, it informs the person that more pills are to be sorted without specifying which pill or which day.

A level 2 assistance is commonly employed when a person first misplaces a pill. The robot is to indicate that there may have been a mistake but is to not provide too much information. For example, if a pill has been placed on Tuesday but that pill is not to be taken on that day, then the selected action may be `lookAndSay('The pills for Tuesday do not look correct.')`.

Level 3. Assistance for level 3 is *verbal directive* assistance, which provides more direct information about how to complete the task. In this case, the robot may indicate not just which pill has been misplaced but also how to fix it. For example, if a green pill has been placed in the afternoon but it should be taken in the morning with breakfast, then the level 3 assistance would be the action `sayText('Green pills are taken with breakfast.')`. Another example is when there are too many pills of a particular color on a given day (e.g., Wednesday). Then the action for a level 3 assistance may be `sayText('You have too many pills on Wednesday.')`.

Level 4. A level 4 assistance adds a *gesture* to provide additional information to the person. The only gesture that the robot employs is to point to a position in the grid. This is done in conjunction with the robot providing verbal assistance. For example, in the scenario in which a green pill has been placed in the afternoon for Tuesday but it should be taken that morning, a level 4 assistance has the robot point at Tuesday on the grid and speak with the action `pointToError(2, 3, 'The green pills on Tuesday should be taken in the morning.')`. The numbers that are the first two parameters represent where the robot is to point. Another example is the case where a blue pill has been placed on a Saturday even though no blue pills are to be taken that day. The action for the level 4 assistance in this case is `pointToError(6, 3, 'The blue pills on Saturday are not needed.')`.

There are two ways for the Medication Management and Assistance component to determine which of these levels of assistance is appropriate. For some classifications of the state of the task, it is clear which level of assistance is necessary. For example, for the “Good” category, there is no reason to provide any assistance above level 1.

In other cases, the component has to keep track of an estimate of how much “need” the person has to determine which assistance level is best. Need is based on how much assistance has been provided during the task. In other words, as more assistance is provided, it becomes apparent that the person has a greater need for assistance. Since the person has more need for assistance, a greater level of assistance should be selected.

The system keeps track of *need* as a number so that it may easily be compared with each level of assistance. The level of assistance immediately greater than the need is typically selected. This is because it is assumed that previous assistance has not been enough, there is now a greater need, and thus the next level of assistance is best. For example, if the *need* currently is 2.4 then an action that provides a level 3 assistance would be selected. More concretely, let us assume that the state of the sorting task is categorized as *Wrong Day*, because there is a blue pill on Saturday that is not

needed. If the current *need* is 2.4, then a level 3 assistance (i.e., `sayText('Try removing the blue pills from Saturday.')`) is selected.

Step 4. Need Update. The final step is to update the estimate of how much need the person has. When the action selected has an assistance level greater than the need, then the need is increased. This allows the robot to provide increasing assistance on each subsequent problem the person is having. When the person starts making progress on the medication sorting, an action with a lower level of assistance is selected, and the estimate of how much need the person has is lowered.

In testing, it was determined that typically the increase in *need* had to be more dramatic than the decrease. Best performance was found when increasing the *need* to equal the level of assistance if the level of assistance is greater than the *need*. If the level of assistance is less than the *need*, then the need is decreased by 0.1 for every successful action by the user. This is best explained with a short example.

Assume the *need* is 1.5 when the person has placed a blue pill on Saturday that should not be taken on Saturday. Initially, the robot will choose an action that has a level 2 assistance (`sayText('Make sure everything is on the right day.')`). At this point the *need* is increased to 2.0. Then the person moves the blue pill to Friday, but blue pills are also not to be taken on Friday. This time the robot will select an action for a level 3 assistance (`'Try removing the blue pills from Friday.'`). Now, the person removes the blue pill and the robot selects an action with a level 1 assistance to notify the person that the action was correct (`sayText('Ok')`). Then the *need* decreases to 2.9. The *need* is not decreased more, because there is recent evidence that the person is having some difficulty. With repeated successes, the *need* will continue to decrease.

4 EVALUATIONS

We describe here a pair of experiments to evaluate the experience of being assisted by a social robot while sorting medications. In the first experiment, university students were recruited to compare the evaluations of a physical robot and a virtual robot. In the second experiment, people with PD were recruited to evaluate a physical robot. The first experiment demonstrated that the robot performed sufficiently well in assisting in sorting medications. Based on the positive evaluations in the first experiment, a follow-up experiment was conducted with the target population. We expected that the people with PD would similarly have positive reactions to the robot. However, results showed that the people with PD generally had a negative reaction. Interviews with the people with PD revealed some of the reasons for their negative evaluations of the robot and its assistance.

4.1 Student Evaluations

In this experiment, university students evaluated either a physical robot or a virtual robot. We expected the physical robot to be preferred by the participants, because previous studies have shown that a physically present robot is perceived more positively and leads to better user performance (see [28] for review). Further details of this experiment can be found in [60] and are summarized here.

Participants. Students ($N = 20$) from an American university were recruited by word of mouth, signs, and web postings to participate in a social robot study. Half were undergraduate students, and the other half were graduate students. Half of the participants were female. Each participant was randomly assigned to one of two conditions: physical robot or virtual robot. Participants were compensated \$10 for their participation.

Materials. A Nao robot was used for both the physical and virtual conditions. The virtual robot was a simulation of the Nao displayed on a screen using Choregraphe. As shown in Figure 5, the

robot was positioned on the table across from the participant, and a medication sorting grid and tray were on the table between the robot and the participant. Also on the table were two pill bottles and a microphone.

The containers of simulated medicines were actual pill bottles. On each bottle, we affixed a label that named that medication and provided instructions on how the medication was to be taken.

After completing the medication sorting, each participant completed a questionnaire consisting of 19 Likert-scale questions, one open question for comments, three demographic questions (major, undergrad/grad, gender), and eight questions about experience with robots. The Likert-scale questions are on a five-point scale ranging from “strongly disagree” to “strongly agree.” The content of the questions covers topics on robot functionality, trust, physical and emotional support, rapport building, safety, social context, and emotional experience. Many questions were adapted from material in other sources: trust [31], physical and emotional support [18], rapport [6], and emotion [48]. The list of questions is as follows:

1. The robot was able to provide you with assistance in the task.
2. The assistance the robot provided was correct.
3. I was able to complete the task more efficiently with the assistance of the robot.
4. When the robot corrected me, I felt included to follow its instructions.
5. I trusted the robot to (correctly) provide assistance.
6. I expected the robot to act in a consistent and predictable manner.
7. The robot was able to provide physical support.
8. The robot was able to provide emotional support.
9. The robot paid attention to me.
10. The robot used action and words that made sense to me.
11. The robot helped me understand how to complete the task.
12. The robot acted in a manner that ensured my safety.
13. The robot was able to warn me of potentially unsafe medication administration.
14. My family would approve of the way the robot assisted me.
15. My care providers would approve of the way the robot assisted me.
16. I felt pleasant during the task.
17. I felt in control of what was happening during the task.
18. I felt I understood what was happening during the task.
19. I felt responsible for completing the task.

Procedure. When the participant entered the room, the robot was on the table in a crouched position with its head down. Before the task started, the research assistant oriented the participant to the task by reading a script. Included in the script were details about the person’s schedule that would affect when the medications were to be taken. The participant was told that the schedule could not be written down. The intent was to create a small load on working memory, hopefully creating enough of a challenge that participants would naturally make mistakes and need assistance from the robot. Additionally, this load on working memory could partially simulate the impact on cognitive function that would be experienced by older adults that could eventually be using our system.

Then the research assistant left the room, and the robot began assisting the participant in sorting the medications. After the medication sorting, the participant completed the questionnaire and then was paid for his or her participation.

4.1.1 Results. Four participants were not able to complete the task due to technical difficulties. To understand the full range of reactions to the robot, data from all participants were included in

Table 1. Results for Each Question for Students

		Physical		Virtual		All		t-test (p)
		M	SD	M	SD	M	SD	P/V
Robot assists	1	4.1	1.29	4	1.05	4.05	1.15	0.426
Robot correct	2	4.1	1.37	4.5	0.71	4.30	1.08	0.211
More efficient	3	4.0	1.15	3.9	1.20	3.95	1.15	0.423
Feel included	4	4.2	0.79	3.7	0.82	3.95	0.83	0.091
Trust robot	5	3.8	1.40	3.9	1.20	3.85	1.27	0.432
Robot consistent	6	3.9	1.20	4.3	1.06	4.10	1.12	0.220
Physical support	7	2.0	1.15	1.7	0.95	1.85	1.04	0.267
Emotional support	8	2.6	1.17	3	0.94	2.80	1.06	0.206
Paid attention	9	2.9	1.37	3	0.82	2.95	1.10	0.423
Made sense	10	4.0	1.25	4.4	0.84	4.20	1.06	0.206
Helped understand	11	3.0	1.25	3.4	1.07	3.20	1.15	0.226
Ensured safety	12	4.0	0.82	4.1	0.74	4.05	0.76	0.389
Able to warn	13	3.4	1.17	3.8	0.92	3.60	1.05	0.204
Family approval	14	3.7	1.34	3.6	1.07	3.65	1.18	0.428
Care providers approval	15	3.7	0.82	3.7	0.95	3.70	0.86	0.500
Felt pleasant	16	3.7	1.34	3.8	0.92	3.75	1.12	0.424
Felt in control	17	4.2	0.92	3.5	1.08	3.85	1.04	0.068
Felt understanding	18	3.9	1.45	3.4	1.35	3.65	1.39	0.218
Felt responsible	19	4.3	1.25	4.2	0.92	4.25	1.07	0.420
	Robot	3.8	0.90	4.0	0.69	3.9	0.78	0.340
	Self	3.7	0.69	3.5	0.76	3.6	0.71	0.298
	Alliance	3.6	1.01	3.6	0.66	3.62	0.83	0.459

the analysis presented here. The means and standard deviations of each question are presented in Table 1. We also used aggregated measures based on a previous principal component analysis, where three components were identified [61]. The “Robot” component measured the robot and its effectiveness, trustworthiness, and reliability. The “Self” component measured how the robot’s assistance affects the person being assisted. The “Alliance” component measured how well the robot supports and includes the person being assisted.

We compared the questionnaires answers of those who interacted with the physical robot with those who interacted with the virtual robot, expecting those who interacted with the physical robot to provide higher ratings. We attempted to verify this using an unpaired t-test, but instead we found very little and no systematic differences between the evaluations of the physical robot and the virtual robot. All comparisons (shown in Table 1 in the P/V column) had $p > 0.05$.

Our hypothesis was that the physical robot would be preferred, but the results did not allow us to reject the null hypothesis. While we could not use these results to assess whether the physical or virtual robot is preferred, we did find that both robots were generally well received. Given the positive evaluations of the robots, we had grounds for furthering the development of the robot by having the target population evaluate the robot.

4.2 People with PD Evaluations

In the follow-up experiment, we sought to continue the development of the robot by evaluating how people with PD respond to its assistance. To get further insights into how people assess the robot, each participant was interviewed after interacting with the robot.

Participants. People with PD ($N = 10$) were recruited by phone to participate in a social robot study. All participants were in early-to-moderate stages of disease severity and had the capacity to independently consent to and participate in the cognitive and physical requirements of the task, with no indication of disoriented mental status. All participants were recruited from a cohort study on community-living adults with Parkinson's disease and had been screened for idiopathic Parkinson's disease by a neurologist and neurological nurse. In addition to the person with PD, each person was invited to bring a care-partner, a person (often a spouse) who assists in the care of the person with PD. Of the 10 participants, three were accompanied by a care-partner. Each participant was paid \$25 for participating in the study.

Materials. A similar setup to the first experiment was used here, but this study focused on only the physical robot to continue with the original motivation in exploring how a physical robot may be able to assist an older adult. One change in the physical setup was that we replaced the sorting tray with a bowl to avoid pills spilling over the edge of the tray.

Instead of the schedule of events for the week being read to the participant by a researcher, each participant was provided with a written schedule of events to be used while sorting medications. The intent was to reduce the cognitive load on the participants and allow the person to focus on completing the task with the robot's assistance. The schedule consisted of a dance class for people with PD on Sunday afternoon and a physical therapy appointment on Tuesday and Thursday mornings. The schedule was displayed in a grid layout similar to the sorting grid and as a list of events for the week.

Procedure. The experiment followed similar procedures as the first, but there were three important differences. First, the participant may have an accompanying care-partner, who watched the whole experiment. The care-partner was invited to watch from the control room with the research assistants but was also welcome to stay in the room with the participant if he or she so chose. If the care-partner stayed in the room, then it was requested that he or she let the participant do the task as much as possible.

To avoid any confusion and to minimize any anxiety around the task, a research assistant introduced the task and did a practice round first. In the practice round, the research assistant instructed the participant to sort one of the medications on the medication sorting grid. The medication being sorted was not one of the medications that the robot would later assist with. After the practice round, the research assistant left the room and the robot continued the experiment as was done in the first experiment.

One small modification made to the robot, compared to the robot used with student participants, was to compensate for the robot not consistently detecting the change in the direction of the gaze of the participant. A manual override was added allowing the researcher in the control room to provide this stimulus in the event that the functionality was failing to operate sufficiently.

Once the task was completed and the participant completed the questionnaire, a research assistant conducted an interview with the person with PD and his or her care-partner. In this interview, we asked the participant and the care-partner about how they manage medications, their expectations and impression of the robot, whether the robot helped too little or too much, and how they would see a robot fitting into their home.

4.2.1 Quantitative Results. Two participants were not able to complete the task due to technical difficulties. As in the previous evaluation, data from all participants were included in the analysis presented here. The results, shown in Table 2, suggest that the people with PD have a generally low evaluation of the interaction with the robot.² Mean of all the scores was 2.60 ($SD = 1.08$),

²We do not provide a statistical comparison with the student evaluations due to subtle differences between the setups.

Table 2. Results for Each Question for People with PD

		M	SD			M	SD
Robot assists	1	2.6	1.26	Helped understand	11	2.5	1.08
Robot correct	2	2.4	1.07	Ensure safety	12	2.5	0.85
More efficient	3	2.2	0.92	Able to warn	13	2.8	1.03
Feel included	4	2.7	1.16	Family approval	14	2.6	0.97
Trust robot	5	2.7	0.82	Care provider approval	15	2.5	0.97
Robot consistent	6	3.3	1.25	Felt pleasant	16	2.7	0.95
Physical support	7	2.1	0.99	Felt in control	17	2.5	1.08
Emotional support	8	1.9	0.88	Felt understanding	18	2.5	1.08
Paid attention	9	2.3	1.25	Felt responsible	19	3.8	1.14
Made sense	10	2.9	1.10	Total	2.6	1.08	

representing a fairly low overall evaluation. Student participants generally had a more positive evaluation, $M = 3.67$, $SD = 1.21$.

The questions with the most negative responses were 2, 3, 7, 8, and 9, all of which had means below 2.5. A low score for question 7 was expected, since the robot had little ability to provide physical assistance. Low scores for 8 and 9 were probably related, as perceptions that the robot is not paying attention to the person can contribute to the robot not providing emotional support.

Questions 2 and 3 indicate that the people with PD felt that the robot's assistance was not correct or helpful. The people that rated questions 2 and 3 lowly may have needed the robot's assistance, but the robot failed to adequately assist. Of the five participants who gave the lowest scores for questions 2 and 3 (mean of 2 or less), three of them were the only ones who reported not having seen a robot in real-life. This suggests that lack of familiarity can contribute to low score for questions 2 and 3.

4.2.2 Qualitative Results. After reading the open answers on the questionnaire and all of the transcripts, we conducted a directed content analysis [20] by developing a coding scheme to identify data related to (1) validity of the experimental setup, (2) performance of the robot, and (3) other technologies. Each topic was further labeled to identify more specific themes. For validity, we coded topics related to comparisons with their current procedures and how they envision the robot fitting in their home. For performance of the robot, we coded topics related to how well or how much the robot assisted and technical issues with the robots performance. Other technologies were further categorized as robot and non-robot. From this analysis, we have found that participants generally expressed that the experiment did not align with their current procedures, there was a lack of assistance provided by the robot, and a robot is unnecessary and other alternatives may be preferable.

Validity of the experimental setup. All of the participants have different procedures and strategies for managing their medications. Seven of the 10 participants described using some sort of pill box, container, or dispenser (and one other used paper cups), and most of the boxes were multi-day containers, similar to our experiment, but not necessarily subdivided into times of day. One participant described some of the similarities and differences as follows:

"I do my medications in a pill box, which has only one, rather than four divisions. It's got the days of the week, and I fill them up with my pharmaceuticals every so often when they're used up, it takes me about 15 to 20 minutes for about 13 pills that I use every day."

In contrast to our setup, the participants generally did not alter their medication schedule to accommodate other life events (e.g., appointments, dinner with friends). As a result, our paradigm, where pills need to be sorted in relation to a schedule of events, caused confusion. Participants described the initial confusion as follows:

- “I didn’t understand at first, the schedule, that I was supposed to ...load pills on this in reference to the schedule.”
- “It wasn’t clear to me how this sheet, how I was supposed to put pills on this in order to just.”
- “I was also struggling cause I wasn’t connecting the label and the schedule.”

While some were confused by this procedure, the robot generally was helpful in this scenario. Participants comments that they were able to get past the initial confusion with the assistance of the robot.

Performance of the robot. While the robot was helpful in this one case, the participants had many concerns about the robot and the potential of it being a part of their regular routine.

Participants comments on the physical performance of the robot:

- “It goes too slow.”
- “The gesturing doesn’t really help that much.”

They also had concerns about the size of the robot:

- “It’s too big.”
- “I think that the robot is too big, I mean you’re not going to carry it around.”
- “It would take up room. I live in a sort of small apartment with my husband and I just can’t imagine how, where the robot would fit in my apartment, it would look really stupid there.”

Some participants were worried about having to program the robot:

- “When the day comes when you have to program it what are you going to do?”
- “You get a medication change, and then he’s obsolete. Then you have to program the computer, the robot, to have the accurate, up-to-date information.”

The robot’s speech also raised concerns:

- “The robot would not repeat what it said.”
- “He doesn’t say much, but when he says it he says it fast.”
- “...because he seemed to talk fast and I couldn’t quite follow what he was saying.”
- “He doesn’t speak clearly.”

In addition to the robot’s speech being a cause for concern, the robot’s poor performance was linked to the fact that the the robot did not always provide sufficient assistance. Seven of the 10 participants wished that the robot provided more instruction and assistance, with one participant expressing frustration over the lack of assistance. For example, in regards to the lack of assistance, participants made the following comments:

- “Something wasn’t right, but it wasn’t telling me, it wasn’t giving me enough information.”
- “I don’t think [it] gave me any assistance.”
- “...maybe not quite enough.”
- “It was inadequate.”

The robot's assistance was not sufficient, and this led to some emotional reactions about the robot. One said the following in regards to the robot's assistance:

"The robot made me feel very frustrated."

In addition to frustration, two participants mentioned being annoyed by the robot. Overall, the lack of assistance appeared to be related to negative emotions.

Other technologies. Since the robot's procedures did not match their own and was too big and did not give sufficient assistance, participants felt the robot was unnecessary and proposed alternatives. Eight of the participants referenced alarms, timers, phones, computers, apps, or some other technical solution. As for the robot being unnecessary, many felt that it would be better to have a simpler solution such as an app on their smartphone or watch. Participants said,

- "[A robot] does seem to be overkill a little bit."
- "It could be like a smartphone."
- "It's too big, and I think you should make it something you can strap on your wrist."
- "I think a computer would be just as good."

Perhaps one reason they wanted a smartphone app was because they felt that what they really needed was assistance in remembering to take their medications. For example, participants suggested the following:

- "It would help remind the person, and I think it would be a great, a great asset."
- "A little timer and it goes off certain times when you're due for your medicine. That seems more realistic than a robot to me."
- "If this would start alarming and maybe send me a message on my phone."
- "...an alarm on the phone or something."
- "Something you wear around your neck, it's a little timer and it goes off certain times when you're due for your medicine"

Perhaps a more intelligent timer would incorporate the following suggestion from another participant:

"I could definitely use somebody saying 'Ok, it's 20 minutes before you want to eat lunch. Now take it, because it's supposed to be taken 20 minutes before a meal.'"

The suggestions that participants had were grounded in their experience with existing technologies, like timers and smartphones. But most participants did not have experience with robots and were not sure what to expect. The point of reference that some people had was industrial robots:

- "Most of the robots I've dealt with were industrial..."
- "I know that they use robots on assembly lines in car manufacturing plants."

Others expressed not knowing what to expect but still mentioned the most salient robot that came to mind:

- "I wasn't really expecting R2D2."
- "I thought it'd be just a computer. And, a certain program that I would go through to help with some task but I didn't expect R2D2."

In addition to these two references to R2D2, two other participants referenced R2D2 or *Star Wars* in the questionnaire. While they said that they did not know what to expect or explicitly said they

were not expecting R2D2, it is probable that their expectations about the robot were influenced by media representations, especially movies like *Star Wars* with the socially assistive robot, R2D2.

4.3 Discussion

For the most part, the participants with PD rated the experience with the robot relatively low, whereas the students generally provided positive ratings. While we do not provide a direct statistical comparison of the participant pools due to slight differences in the setups of the evaluations, it is clear that the people with PD had a very different experience interacting with the robot than the student participants did. As a consequence, the results obtained from the student participants lacked external validity and provided insufficient feedback regarding the effectiveness of the robot to assist a person with Parkinson's disease.

The differences between the evaluations provided by the students and the people with PD could be the result of differences in the setups, but it is unlikely that small procedural differences (e.g., demonstrating how to do the task) would have had such a great effect. It is more probable that people with PD gave such negative evaluations as the result of some of other causes, including people with PD having different (1) routines for managing medications, (2) needs that technology may assist with, and (3) familiarity with and expectations of technology. Each of the differences can be informative for how to proceed with future designs and evaluations.

1. Medication management procedures. One source of negative evaluations of the robot is that the participants with PD did not see how the robot could integrate into their daily routines, because the procedure for medication sorting that the robot used differed from how the participants manage their medications. Each of the participants with PD in the study has a particular routine used for managing their medications. Some of them use daily pill organizers, and others just use reminders to notify them when it is time to take a dose of a medication. They work with their doctors and care providers to customize a medication regimen and routine that works for them. The robot, however, was designed to follow a specific medication sorting procedure based on an in-depth task analysis [61] of a standardized medication sorting task [40]. While the procedure used by the robot for medication sorting was unfamiliar to both pools of participants, the rigidly defined task does allow us to investigate how a socially assistive robot may provide varying levels of assistance to aid in a task while trying to maintain the autonomy and dignity of the person being assisted. However, the unusual procedure caused confusion for some of the participants with PD confusing to some participants.

Since the people with PD were evaluating the robot from the perspective of a current need, they were able to identify many of the shortcomings of the robot and became frustrated with the robot's failures. Interviews with these participants revealed that they found the robot to be unnecessary and not meeting their current needs. As an alternative, some proposed other (more mature) technologies that may have a more immediate benefit. They suggested alarms and smartphone apps, because these technologies are more familiar and have been demonstrated to be effective in other contexts.

Contrasting the perspectives of the participants with PD and the student participants, most of whom probably do not have a complex medication regimen, the students did not assess the robot for its utility in their daily lives. Instead, they assessed the robot for some hypothetical situation where someone would need assistance. This drastically different perspective makes it difficult for them to recognize deficiencies in how the robot operates and likely contributes to their high evaluations.

Unlike the student participants, many older adults establish rigorous routines for managing their medications [14], and future work should use participatory design activities to understand

the varying routines people use and to determine the variety of needs and preferences among individuals as they follow their routines. Working with older adults as co-designers would allow us to learn more about their preferences as they engage in their routines and enable us to further investigate how to integrate into their existing routines and explore how routines can be customized and personalized to better allow for the robot to address individuals' needs. This process of developing personalized and more realistic scenarios that cover many of the possible variants in routines will require significant resources, but it will achieve greater ecological validity, allow us to be able to more accurately assess the effectiveness of the robot, and ultimately better serve the end users.

The current design of the robot supports only one rigidly defined procedure. To enable the robot to be more flexible in the medication sorting procedures, components of the robot architecture will need to be revised. In particular, the Vision and Medication Management and Assistance components need updates to accommodate different pill containers, the frequency with which individuals organize their medications, and characteristics of particular medications and their prescribed usage. The current Vision component, which is configured to recognize a specific sorting grid and only two types of pills, would need to be enhanced to recognize other containers and types of pills. To provide more flexibility in procedures and constraints, updating the Medication Management and Assistance to use a plan-based approach to selecting assistance [62] would enable the robot to plan out a valid sequence of sorting steps for a set of constraints (prescriptions, vitamins, meals, life events, etc.) it has not seen before (i.e., unique to that individual).

2. Needs of older adults. Another source of negative evaluations from the people with PD is that the robot did not meet their needs. Recognizing that the medication sorting procedure followed by the robot differs from the daily routines of the people with PD, we can still try to assess how well the robot assists at this unfamiliar procedure for managing medications. Our assessment indicated that the robot was not able to consistently meet the needs of the people with PD. People with PD may have had different (or simply greater) needs than the student participants, and the robot may have provided *too little* assistance. As evidence for the robot giving too little assistance, the participants reported that the robot was not able to help them complete the task more correctly (question 2) and efficiently (question 3). Additionally, participants described the robot's assistance as "not quite enough" and "inadequate."

If a person is not given enough assistance, then the person may be left not knowing what to do. When feeling responsible but unclear how to proceed, a person may feel like he or she is not in control of the task. Not knowing what to do and not feeling in control does not empower the participant to manage his or her own health and instead may diminish the person's sense of autonomy and lead to some of the negative emotions (e.g., frustration and annoyance) expressed in the interviews. Conversely, if the robot provides more assistance (but not too much), then the participant would have a greater ability to complete the task, have greater faith that the robot can help, and not feel incompetent about one's own ability to do the task.

One reason why the robot may have given too little assistance was a design choice to minimize assistance to maximize autonomy. This approach worked well with students, who rarely needed much assistance. With older adults, or specifically those with PD, the minimal assistance did not meet the need of the participant. With the assistance being insufficient, the need the person had grown, creating a gap between the assistance provided and the assistance needed. To remedy this situation, future designs need to better adapt to the needs of the individual, as there is great diversity in the needs of older adults as they age [21]. Methods for improving the adaptation could happen before, during, or after the interaction. Before the interaction, the robot could conduct an interview to assess needs, similar to that done in [6], and customize the following interaction

based on the interview results. Also, while the robot is assisting in medication sorting, it needs to better detect when the person needs assistance and how much assistance to provide assistance that matches those needs. Furthermore, after each interaction with the robot, it could retrospect on the completed interaction to learn the needs of the person so that it may better address those needs in future interactions. To enable these adaptations, the Medication Management and Assistance component would need to be enhanced to take into account these factors during the decision-making process. Additionally, new components may also be necessary to handle the learning aspects.

In addition to the robot not giving enough assistance, another reason the robot was seen as not helpful is that it did not always give the type of assistance that was needed. In particular, some participants were frustrated by the fact that they could not understand the speech of the robot. They critiqued the volume was too low or that it simply was hard to understand, and some would have liked the robot to repeat itself. The speech of the Nao robot can be a little difficult to understand, and for someone with any amount of hearing degradation, it can be even more challenging to understand. These challenges are less likely to occur with a younger population who tend to have better hearing, and thus this is another example of how older adults have different needs that the robot was not able to address.

For the robot to better accommodate the needs of the older adults, the robot should be able to understand direct commands (e.g., repeat what it said or adjust volume, or point at a mistake). This would allow the user to be more in control of the robot and thus more in control of the task. For the robot to understand such commands, the architecture would need some natural language understanding components that go beyond the simple keyword matching that the current design employs. Being able to give commands to the robot could go beyond direct commands and allow the user to indicate needs and preferences. Some older adults have expressed the desire to be able to specify to the robot what they want it to do or not do [8], thereby restricting the robot's autonomy to advance the user's autonomy.

3. Familiarity with other technologies. The evaluations of the participants may have been biased by the technology with which they are more familiar. In the interviews with the participants with PD, most referenced some other technology with which they are more familiar (e.g., alarms, smartphones, computers, other robots). Their perception of the current robot has some foundation in these prior experiences [55], as people reference familiar systems when introduced to new technologies [7]. In addition to referencing other relevant technologies, participants with PD referenced other robots, some of which they had interacted with (e.g., manufacturing robots), others that they had seen in movies. Four of the participants referenced R2D2 or *Star Wars*, possibly suggesting they are anchoring on a robot with far greater capabilities, leading to disappointment in the current robot.

The student participants likely also have had experience with using alarms, reminders, and a variety of other technologies but not necessarily in regards to medications. Whatever bias the students may have based on prior experiences with these technologies could be removed from their consideration, because they are already in the mindset of evaluating the robot for a hypothetical scenario, as we described above.

In future work, we need to recognize that familiarity and experience can influence perceptions and consider how to integrate with existing technologies while also providing training for newer technologies. One benefit of integrating with existing technologies is that it can lead to better technology adoption [24], hopefully leading to the users reaping the benefits of the adopted technology. Since many of the participants with PD referenced technologies that can be used to remind them to take their medications, future solutions can leverage this familiarity by linking the sorting task to the reminding functionality of known technologies. For example, we could create a natural

transition from medication sorting into reminding by allowing the completed sorting to be used as the basis of reminders. These reminders do not need to originate from the robot, as it is unreasonable to have the robot travel with the person (due to size, mostly), but instead the robot can push the reminders to a calendaring system that can provide alerts to a mobile device, such as a smartphone.

Integrating with existing technologies would better allow participants to envision how the system could benefit them in their lives, but integrating these technologies into their routines may still be challenging. We need to continue to work with OTs and other healthcare providers to develop strategies to train people to work with the robot and to help modify routines to integrate the robot.

OTs can be a valuable resource in helping to design a robot to meet the needs of people with PD. They have knowledge and experience with the daily activities that their clients need to be able to perform to have a high quality of life. For this reason, we have relied on OTs to provide expert knowledge of how to assist people with PD, including the development of our initial task analysis and guidance in the robot architecture [61]. However, it may be insufficient to solely rely on these experts, and we need to do more to capture the perspectives of the older adults that would be using the robot. To accomplish this, a participatory design approach, where the older adults are co-designers of the robot, would allow the researchers and the older adults to learn from each other, working together to design a solution that is technologically feasible while also grounded in the realities of their domestic environment [1]. As mentioned above, co-designing would allow us to design a robot that better integrates with the routines of older adults. Additionally, co-designing with older adults would allow us to further investigate human autonomy in the context of robot control. A previous participatory design project investigated the challenges in balancing the need to control the behavior of the robot, which may lead to more user autonomy, with allowing the robot to autonomously provide the necessary assistance [8].

This is just one example of the unique needs and perspectives of older adults that make it imperative that they are involved in the design and evaluation of robots designed to assist them. However, including some older adults, such as ones with Parkinson's disease, can be challenging due to limited access to a rare, vulnerable population [30]. This inherent limitation has the potential to impede progress on designing robots for this population, but new experimental methods could allow for increased participation. For example, we should consider using telepresence experiments or "expert subjects," who are trained advocates for a given population. If we do not consider alternative approaches to increasing participation of people with PD, then the likely result will be fully autonomous social robots not being available to this population in the near term.

5 FUTURE WORK

To bring us closer to realizing a social robot to assist in medication sorting, we must work with the potential end users to better understand the ecology of the task. One theme that has become evident in the work thus far is that the robot must better adapt to the individual. There are at least three ways in which the robot should adapt, including (1) better adapting to an individual's level of need, (2) adapting to individualized procedures to accomplish medication sorting, and (3) customizing to an individual's medications and life events.

Improving the estimation of a person's need would enable the robot to provide assistance that better matches the person's needs and expectations. In the system described here, the robot was conservative in estimating need, often underestimating the amount of need. Ongoing work is looking at improved mechanisms for updating the estimate of need, incorporating a measure of the person's progress in completing the task [62].

To address issues regarding the rigidity of the procedures for completing a medication sorting task, the robot needs to be able to adapt to an individual's procedure, possibly adapting as the task progresses. One person may prefer to organize one medication at a time, sorting all of the pills for that medication before moving to the next one. Meanwhile, another person may prefer to organize things one day at a time, sorting all the medications for a given day before progressing to the next day. One way for the robot to adapt to these different strategies is for the robot to incorporate a planner into its decision-making capabilities, allowing the robot to shift to a different sequence of steps that accomplish the same goal [62]. Once a person's preferred strategies are identified, after repeated interactions, the robot can begin to assume how the person would like to complete the task and give assistance in accordance with that preference.

In the procedures we used in our evaluations, we defined a pair of medications to be sorted and some life events that the person needed to consider when sorting the medications. The medications and life events should be defined by the person and not by the robot. One possibility is for the robot to interface with another system that provides this information. Alternatively, a solution that keeps the person involved in the sharing of information would be for the robot to ask the person about the medications and any life events. As a result of the robot's inquiries, it can modify constraints defining a valid goal. For example, one person may be visiting family on Monday and requires taking some medications before leaving the house that day. The robot would incorporate this information and guide the person to a successful sorting of the medications that accommodates the person's travel plans.

6 CONCLUSION

There are numerous challenges in deploying a fully autonomous, socially assistive robot, and there are even more challenges for robots designed to assist older adults in complex tasks and routines that they must accomplish to maintain their autonomy in home and community. In addition to the necessary robustness of the system, the design of social robots for older adults must recognize that older adults have different needs (and typically greater health care needs) than younger populations and that older adults' expectations and familiarity with technology often differs from younger populations. Furthermore, when targeting a population that is rare and difficult to recruit for experiments, because they live with a high-burden disease that restricts their time or ability to participate, there is a clear challenge of including them early in the design process, which, however, is important to ensure that the research has the external validity necessary to make a real impact. All of these challenges, which are not insurmountable, do present barriers to providing effective solutions to the target populations, and it becomes unlikely that fully autonomous socially assistive robots to assist older adults will be available in the near term.

While we may not be able to deploy these robots soon, there are many important steps to take and important milestones along the way. For example, it is important to gain a better understanding of how people with PD react to a socially assistive robot, how it makes them feel, and how it affects their autonomy. This will be critical to developing an assistive technology that will be accepted by the target population and, therefore, be able to make a desirable impact on their quality of life. It is clear from our own experience that only using university student subjects during the development and evaluation phases is insufficient and might give a skewed picture of the system's appropriateness and effectiveness compared to the target population. At the same time, as it is impossible to recruit a significant number of subjects of the target population for ongoing studies, it is important for us to devise new experimental methods that allow for early participation of the target population in the design process, e.g., using telepresence experiments or "expert subjects" (who are eager to help the development of assistive robots and are willing to come to the lab more frequently for pilot studies, where they can provide explicit feedback on the design of the system).

ACKNOWLEDGMENTS

The authors thank the many research assistants that contributed to this work: Nah Young Lee, Annie Saechao, Sharon Hershenson, Luca Brusatin, and Evan Krause.

REFERENCES

- [1] H. R. Lee, Šabanović, W. L. Chang, S. Nagata, J. Piatt, C. Bennett, and D. Hakken. 2017. Steps toward participatory design of social robots: mutual learning with older adults with depression. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. 244–253.
- [2] Ronald C. Arkin, Matthias Scheutz, and Linda Tickle-Degnen. 2014. Preserving dignity in patient caregiver relationships using moral emotions and robots. In *Proceedings of the IEEE International Symposium on Ethics in Science, Technology and Engineering*. 1–5.
- [3] American Occupational Therapy Association et al. 2015. Occupational therapy code of ethics (2015). *Amer. J. Occup. Ther.* 69, 3 (2015). <https://ajot.aota.org/article.aspx?articleid=2442685>.
- [4] Jacquelyn L. Bainbridge and J. Mark Ruscin. 2009. Challenges of treatment adherence in older patients with Parkinson's disease. *Drugs Aging* 26, 2 (2009), 145–155.
- [5] Elliot Baum. 2006. Pill dispenser. U.S. Patent 7,100,793.
- [6] Priscilla Briggs, Matthias Scheutz, and Linda Tickle-deggen. 2015. Are robots ready for administering health status surveys? First results from an HRI study with subjects with Parkinson's disease. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 327–334.
- [7] Susan A. Brown and Viswanath Venkatesh. 2005. Model of adoption of technology in households: A baseline model test and extension incorporating household life cycle. *MIS Quart.* 29, 3 (2005).
- [8] Praminda Caleb-Solly, Sanja Dogramadzi, David Ellender, Tina Fear, and Herjan van den Heuvel. 2014. A mixed-method approach to evoke creative and holistic thinking about robots in a home environment. In *Proceedings of the ACM/IEEE International Conference on Human-robot Interaction*. ACM, 374–381.
- [9] Vicki S. Conn, Adam R. Haf Dahl, Pamela S. Cooper, Todd M. Rupp, David R. Mehr, and Cynthia L. Russell. 2009. Interventions to improve medication adherence among older adults: Meta-analysis of adherence outcomes among randomized controlled trials. *Gerontologist* 49, 4 (2009), 447–462.
- [10] Chandan Datta, H. Yul Yang, Priyesh Tiwari, I. Han Kuo, and Bruce A. MacDonald. 2011. End user programming to enable closed-loop medication management using a healthcare robot. In *Proceedings of Australasian Conference on Robotics and Automation*.
- [11] Lindsey Dayer, Seth Heldenbrand, Paul Anderson, Paul O. Gubbins, and Bradley C. Martin. 2013. Smartphone medication adherence apps: Potential benefits to patients and providers. *J. Amer. Pharm. Assoc.* 53, 2 (2013), 172–181.
- [12] E. R. Foster. 2014. Instrumental activities of daily living performance among people with Parkinson's disease without dementia. *Amer. J. Occup. Ther.* 68, 3 (2014), 353–362.
- [13] Donald Grosset, Angelo Antonini, Margherita Canesi, Gianni Pezzoli, Andrew Lees, Karen Shaw, Esther Cubo, Pablo Martinez-Martin, Olivier Rascol, Laurence Negre-Pages, et al. 2009. Adherence to antiparkinson medication in a multicenter european study. *Movement Disorders* 24, 6 (2009), 826–832.
- [14] Jörg W. Haslbeck and Doris Schaeffer. 2009. Routines in medication management: The perspective of people with chronic conditions. *Chronic Illness* 5, 3 (2009), 184–196.
- [15] Leah M. Haverhals, Courtney A. Lee, Katie A. Siek, Carol A. Darr, Sunny A. Linnebur, J. Mark Ruscin, and Stephen E. Ross. 2011. Older adults with multi-morbidity: Medication management processes and design implications for personal health applications. *J. Med. Internet Res.* 13, 2 (2011), e44.
- [16] Tamara L. Hayes, John M. Hunt, Andre Adami, and Jeffrey A. Kaye. 2006. An electronic pillbox for continuous monitoring of medication adherence. In *Proceedings of the International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 6400–6403.
- [17] Timothy J. Hazen and Issam Bazzi. 2001. A comparison and combination of methods for OOV word detection and word confidence scoring. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'01)*, Vol. 1. IEEE, 397–400.
- [18] Amanda R. Hemmesch. 2011. *The Stigmatizing Effects of Facial Masking and Abnormal Bodily Movement on Older Adults' First Impressions of Individuals with Parkinson's Disease*. Brandeis University.
- [19] Judith E. Hertz and Cynthia A. Anschutz. 2002. Relationships among perceived enactment of autonomy, self-care, and holistic health in community-dwelling older adults. *J. Holistic Nurs.* 20, 2 (2002), 166–186.
- [20] Hsiu-Fang Hsieh and Sarah E. Shannon. 2005. Three approaches to qualitative content analysis. *Qualitat. Health Res.* 15, 9 (2005), 1277–1288.
- [21] Almar A. L. Kok, Marja J. Aartsen, Dorly J. H. Deeg, and Martijn Huisman. 2015. Capturing the diversity of successful aging: An operational definition based on 16-year trajectories of functioning. *Gerontologist* 57, 2 (2015), 240–251.

- [22] Amit S. Kulkarni, Rajesh Balkrishnan, Roger T. Anderson, Heather M. Edin, Jeff Kirsch, and Mark A. Stacy. 2008. Medication adherence and associated outcomes in medicare health maintenance organization-enrolled older adults with Parkinson's disease. *Movement Disorders* 23, 3 (2008), 359–365.
- [23] Paul Lamere, Philip Kwok, Evandro Gouvea, Bhiksha Raj, Rita Singh, William Walker, Manfred Warmuth, and Peter Wolf. 2003. The CMU SPHINX-4 speech recognition system. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'03)*, Vol. 1. Citeseer, 2–5.
- [24] Chaiwoo Lee and Joseph F. Coughlin. 2015. PERSPECTIVE: Older adults' adoption of technology: An integrated approach to identifying determinants and barriers. *J. Product Innovat. Manage.* 32, 5 (2015), 747–759.
- [25] Hee Rin Lee and Laurel D. Riek. 2018. Reframing assistive robots to promote successful aging. *ACM Trans. Hum.-Robot Interact.* 7, 1 (2018), 11.
- [26] Matthew L. Lee and Anind K. Dey. 2011. Reflecting on pills and phone use: Supporting awareness of functional abilities for older adults. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2095–2104.
- [27] Wendy Levinson, Audiey Kao, Alma Kuby, and Ronald A. Thisted. 2005. Not all patients want to participate in decision making: A national study of public preferences. *J. Gen. Internal Med.* 20, 6 (2005), 531–535.
- [28] Jany Li. 2015. The benefit of being physically present: A survey of experimental works comparing copresent robots, telepresent robots and virtual agents. *Int. J. Hum. Comput. Studies* 77 (Jan. 2015), 23–27.
- [29] Eric J. MacLaughlin, Cynthia L. Raehl, Angela K. Treadway, Teresa L. Sterling, Dennis P. Zoller, and Chester A. Bond. 2005. Assessing medication adherence in the elderly: Which tools to use in clinical practice? *Drugs Aging* 22, 3 (2005), 231–255.
- [30] C. Marras, J. C. Beck, J. H. Bower, E. Roberts, B. Ritz, G. W. Ross, R. D. Abbott, R. Savica, S. K. Van Den Eeden, A. W. Willis et al. 2018. Prevalence of Parkinson's disease across North America. *NPJ Parkinson's Dis.* 4, 1 (2018), 21.
- [31] Bonnie M. Muir and Neville Moray. 1996. Trust in automation. Part II. Experimental studies of trust and human intervention in a process control simulation. *Ergonomics* 39, 3 (1996), 429–460.
- [32] Alan Naditz. 2008. Medication compliance—helping patients through technology: Modern “smart” pillboxes keep memory-short patients on their medical regimen. *Telemed. e-Health* 14, 9 (2008), 875–880.
- [33] Lennart Nordenfelt. 2004. The varieties of dignity. *Health Care Anal.* 12, 2 (2004), 69–81.
- [34] Lars Osterberg and Terrence Blaschke. 2005. Adherence to medication. *New Engl. J. Med.* 353, 5 (2005), 487–497.
- [35] T. J. Parker. 2014. Prescribing design. *Design Manage. Rev.* 25, 3 (2014), 30–33.
- [36] Jürgen Pripfl, Tobias Körtner, Daliah Batko-Klein, Denise Hebesberger, Markus Weninger, Christoph Gisinger, Susanne Frennert, Hakan Efring, Margarita Antona, Ilia Adami, et al. 2016. Results of a real world trial with a mobile social service robot for older adults. In *Proceedings of the 11th ACM/IEEE International Conference on Human Robot Interaction*. IEEE Press, 497–498.
- [37] Ireen M. Proot, Harry F. J. M. Crebolder, Huda Huijer Abu-Saad, Ton H. G. M. Macor, and Ruud H. J. Ter Meulen. 2000. Facilitating and constraining factors on autonomy: The views of stroke patients on admission into nursing homes. *Clin. Nurs. Res.* 9, 4 (2000), 460–478.
- [38] Islam Qudah, Peter Leijdekkers, and Valerie Gay. 2010. Using mobile phones to improve medication compliance and awareness for cardiac patients. In *Proceedings of the 3rd International Conference on Pervasive Technologies Related to Assistive Environments*. ACM, 36.
- [39] M. V. Radomski and C. A. Trombly (Eds.). 2013. *Occupational Therapy for Physical Dysfunction* (7th ed.). Lippincott Williams & Wilkins.
- [40] J. C. Rogers and M. B. Holm. 1994. Performance Assessment of Self-Care Skills (PASS-Home) Version 3.1 [Unpublished assessment tool]. University of Pittsburgh, Pittsburgh, PA.
- [41] Richard M. Ryan, Jennifer G. La Guardia, Jessica Solky-Butzel, Valery Chirkov, and Youngmee Kim. 2005. On the interpersonal regulation of emotions: Emotional reliance across gender, relationships, and cultures. *Personal Relation.* 12, 1 (2005), 145–163.
- [42] Anil Sahai, Stephen K. Breede, Roger O. Topliffe, and Douglas A. Topliffe. 1999. Medication dispensing system. U.S. Patent 5,971,594.
- [43] Matthias Scheutz. 2006. ADE—Steps towards a distributed development and runtime environment for complex robotic agent architectures. *Appl. Artif. Intell.* 20, 4–5 (2006).
- [44] Matthias Scheutz, Gordon Briggs, Rehj Cantrell, Evan Krause, Tom Williams, and Richard Veale. 2013. Novel mechanisms for natural human-robot interactions in the DIARC architecture. In *Proceedings of the AAAI Workshop on Intelligent Robotic Systems*.
- [45] Ju Young Shin, Barbara Habermann, and Ingrid Pretzer-Aboff. 2015. Challenges and strategies of medication adherence in Parkinson's disease: A qualitative study. *Geriatr. Nurs.* 36, 3 (2015), 192–196.
- [46] Candace L. Sidner, Christopher Lee, Cory D. Kidd, Neal Lesh, and Charles Rich. 2005. Explorations in engagement for humans and robots. *Artific. Intell.* 166, 1–2 (Aug. 2005), 140–164.

- [47] Cory-ann Smarr, Akanksha Prakash, Jenay M. Beer, Tracy L. Mitzner, Charles C. Kemp, and Wendy A. Rogers. 2012. Older adults' preferences for and acceptance of robot assistance for everyday living tasks. In *Proceedings of the Human Factors and Ergonomics Society 56th Annual Meeting*. 153–157.
- [48] Craig A. Smith and Phoebe C. Ellsworth. 1985. Patterns of cognitive appraisal in emotion. *J. Personal. Soc. Psychol.* 48, 4 (1985), 813–838.
- [49] Michael A. Steinman and Joseph T. Hanlon. 2010. Managing medications in clinically complex elders: "There's got to be a happy medium." *J. Amer. Med. Assoc.* 304, 14 (2010), 1592–1601.
- [50] Carol S. Stillel, Catherine M. Bender, Jacqueline Dunbar-Jacob, Susan Sereika, and Christopher M. Ryan. 2010. The impact of cognitive function on medication management: Three studies. *Health Psychol.* 29, 1 (2010), 50.
- [51] P. Tiwari, J. Warren, K. Day, and C. Datta. 2011. Comprehensive support for self management of medications by a networked robot for the elderly. In *Proceedings of the Health Informatics New Zealand 10th Annual Conference and Exhibition*.
- [52] Martijn van der Eijk, Frouke A. P. Nijhuis, Marjan J. Faber, and Bastiaan R. Bloem. 2013. Moving from physician-centered care towards patient-centered care for Parkinson's disease patients. *Parkinson. Relat. Disord.* 19, 11 (2013), 923–927.
- [53] Kazuyoshi Wada, Takanori Shibata, Tomoko Saito, Kayoko Sakamoto, and Kazuo Tanie. 2005. Psychological and social effects of one year robot assisted activity on elderly people at a health service facility for the aged. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Vol. 2005. 2785–2790.
- [54] R. L. Wade, B. Clancey, and J. Michaeli. 2016. Improvement in antihypertensive and cholesterol-lowering medication persistence using a mobile technology application. *Value Health* 19, 3 (2 2016), A306.
- [55] Kieran Walsh and Aoife Callan. 2011. Perceptions, preferences, and acceptance of information and communication technologies in older-adult community care settings in Ireland: A case-study and ranked-care program analysis. *Ageing Int.* 36, 1 (2011), 102–122.
- [56] Ilene Warner. 1998. Telemedicine in home health care: The current status of practice. *Home Health Care Manage. Pract.* 10, 2 (1998), 62–68.
- [57] Netta Weinstein, Andrew K. Przybylski, and Richard M. Ryan. 2012. The index of autonomous functioning: Development of a scale of human autonomy. *J. Res. Personal.* 46, 4 (2012), 397–413.
- [58] Jason R. Wilson, Nah Young Lee, Annie Saechao, Sharon Hershenson, Matthias Scheutz, and Linda Tickle-Degnen. 2017. Hand gestures and verbal acknowledgments improve human-robot rapport. In *Proceedings of the International Conference on Social Robotics*. Springer, 334–344.
- [59] Jason R. Wilson, Nah Young Lee, Annie Saechao, and Matthias Scheutz. 2016. Autonomy and dignity : Principles in designing effective social robots to assist in the care of older adults. In *Proceedings of the Workshop on Using Social Robots to Improve the Quality of Life in the Elderly (ICSR'16)*.
- [60] Jason R. Wilson, Nah Young Lee, Annie Saechao, Linda Tickle-Degnen, and Matthias Scheutz. 2016. Designing a social robot to assist in medication sorting. In *Social Robotics (ICSR'16)*, A. Agah, J. J. Cabibihan, A. Howard, M. Salichs, and H. He (Eds.). Lecture Notes in Computer Science, vol. 9979. Springer, Cham. https://link.springer.com/chapter/10.1007/978-3-319-47437-3_21.
- [61] Jason R. Wilson, Linda Tickle-Degnen, and Matthias Scheutz. 2016. Designing a social robot to assist in medication sorting. In *Proceedings of the International Conference on Social Robotics*. Springer International Publishing, 211–221.
- [62] Jason R. Wilson, Michael Wransky, and Jorge Tierno. 2018. General approach to automatically generating need-based assistance. In *Proceedings of the 6th Annual Conference on Advances in Cognitive Systems*.

Received December 2018; revised August 2019; accepted January 2020