Cyber-physical Support of Daily Activities: A Robot/Smart Home Partnership

CHRISTOPHER PEREYDA, School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA, USA

NISHA RAGHUNATH, Department of Psychology, Washington State University, Pullman, WA, USA

BRYAN MINOR and GARRETT WILSON, School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA, USA

MAUREEN SCHMITTER-EDGECOMBE, Department of Psychology, Washington State University, Pullman, WA, USA

DIANE J. COOK, School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA, USA

This article introduces RAS, a cyber-physical system that supports individuals with memory limitations to perform daily activities in their own homes. RAS represents a partnership between a smart home, a robot, and software agents. When smart home residents perform activities, RAS senses their movement in the space and identifies the current activity. RAS tracks activity steps to detect omission errors. When an error is detected, the RAS robot finds and approaches the human with an offer of assistance. Assistance consists of playing a video recording of the entire activity, showing the omitted activity step, or guiding the resident to the object that is required for the current step. We evaluated RAS performance for 54 participants performing three scripted activities in a smart home testbed and for 2 participants using the system over multiple days in their own homes. In the testbed experiment, activity errors were detected with a sensitivity of 0.955 and specificity of 0.992. RAS assistance was performed successfully with a rate of 0.600. In the in-home experiments, activity errors were detected with a combined sensitivity of 0.905 and a combined specificity of 0.988. RAS assistance was performed successfully for the in-home experiments with a rate of 0.830.

CCS Concepts: • **Human-centered computing** → *Ubiquitous and mobile computing*; • **Computing methodologies** → *Machine learning*; • **Computer systems organization** → *Embedded and cyber-physical systems*; *Robotics*;

Additional Key Words and Phrases: Smart homes, robotics, activity recognition, activity-aware user prompting

This material is based upon work supported by the National Institutes of Health under grant R25AG046114 and the National Science Foundation under grants 1734558, 1543656, and 1757632.

Authors' addresses: C. Pereyda, B. Minor, G. Wilson, and D. J. Cook, School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA, USA; emails: {christopher.pereyda, bminor, garrett.wilson, djcook}@wsu.edu; N. Raghunath and M. Schmitter-Edgecombe, Department of Psychology, Washington State University, Pullman, WA, USA; emails: {nisha.raghunath, schmitter-e}@wsu.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 ACM 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 Association for Computing Machinery.

2378-962X/2019/12-ART21 \$15.00

<https://doi.org/10.1145/3365225>

ACM Reference format:

Christopher Pereyda, Nisha Raghunath, Bryan Minor, Garrett Wilson, Maureen Schmitter-Edgecombe, and Diane J. Cook. 2019. Cyber-physical Support of Daily Activities: A Robot/Smart Home Partnership. *ACM Trans. Cyber-Phys. Syst.* 4, 2, Article 21 (December 2019), 24 pages.
<https://doi.org/10.1145/3365225>

1 INTRODUCTION

Advances in cyber-physical systems have made machines in our world more responsive, adaptive, precise, and efficient. The impact of these transformed systems has propelled dramatic increases in both CPS research and commercialization [1, 2]. Because of the aging of the population, a high-priority need for cyber-physical support is assisting older adults with daily activities. Health care advances have increased life expectancy, with a result that older people are projected to outnumber children for the first time in U.S. history [3]. The ratio of working-age adults to older adults will fall to just 2.5 by 2060, thus global aging will result in health care needs that cannot be met by family, friends, and care providers. Some of the needs introduced by this age tsunami can be met by technology [4] through automation of health assessment and provision of assistance to both extend functional independence and attenuate the impact of cognitive decline.

Machine learning-driven smart homes deliver valuable health monitoring and assessment technologies. By collecting and analyzing sensor data from wearables, ambient sensors, and cameras, the technologies can monitor activity level, gait, sleep patterns, computer usage, social interactions, and daily behavior patterns [5–14]. This information has been employed to assess depression and loneliness [15–17], rehabilitation [18, 19], schizophrenia and other targeted health conditions [20, 21], fall risk [22–24], and cognitive performance [25–28]. However, these monitoring technologies need to collaborate with other systems that can provide physical interaction with older adults as well as physical assistance when needed.

Our goal is to partner smart home-based sensing technologies with machine learning-based reasoning and robotic assistance to provide a full activity support cycle of sensing, identifying, assessing, and acting to support a person's daily routine. In support of this goal, we introduce RAS, a Robotic Activity Support cyber-physical system. As Figure 1 illustrates, the RAS smart home component uses embedded sensors to monitor individuals as they perform their daily routines. RAS analyzes collected sensor data to identify activities as they occur. Comparing the activity sequence with typical occurrences of the activity, the system determines if an activity step is missing or performed out of order. If such an activity error is detected, then the RAS robot agent acts by approaching the smart home resident and offering support. The RAS robot acts like a butler, offering to play a video of the missing step, play a video of the entire activity, or lead the individual to where the object is located that they need for the activity step.

While RAS does not currently play all of the roles associated with a robot butler, it does play an important role for individuals experiencing memory impairment. Memory impacts the ability of individuals to function safely and independently, particularly as cognition-affecting diseases progress. In such situations, RAS acts as a cognitive prosthesis. By catching errors in daily activities, reminding individuals how to perform those activities, and providing a physical interactive presence, the system can offer activity support that assists individuals in independently performing basic and instrumental activities of daily living (ADLs).

In this article, we introduce the RAS cyber-physical system and its technical elements. We also investigate three issues related to the usability of such a system for real-world activity support. First, we want to determine whether the proposed cyber-physical infrastructure can support robust activity recognition and tracking with a diversity of sensors and smart home residents. We

Fig. 1. RAS system overview. *Sense*: sensors in the smart home collect data while the resident performs daily tasks. *Identify*: activity recognition identifies activities as they occur. *Assess*: RAS detects missing steps in the performed activity. *Act*: the RAS robot approaches the resident to offer help with the activity.

will assess this by recognizing and tracking a set of activities performed by participants in smart home settings, first with item sensors and then with motion and door sensors. Second, we want to assess how usable a robotic activity support CPS would be for our target population, older adults with cognitive impairment. We will assess this issue by gathering and comparing feedback from both younger and older adults who use RAS to monitor and correct activities. Third, we want to determine whether a complex cyber physical system such as RAS can be used in natural, everyday settings. To assess this issue, we collect feedback from two sets of participants: one set performed scripted activities in a testbed environment and a second set performed daily activities in their own homes.

2 RELATED WORK

RAS marks a next step in the development of medical cyber-physical systems. As recent surveys describe, healthcare cyber-physical systems represent a critical integration of medical devices [29, 30]. Medical CPS devices include hardware such as ambient and wearable sensors, data components such as cloud-based storage/processing and integrated electronic health records, and interface components to support remote caregiver communication. Related research has largely focused on integration challenges including component inter-operability, security/privacy, and quality of service. Some previous work, however, has introduced cyber-physical systems that integrate both sensing and robot assistance. Specifically, Do et al. [31] design a robot-integrated smart home, RiSH, to detect and respond to human falls. While residents wear sensing devices to monitor potential falls, fall detection is further enhanced by fusing wearable accelerometer data with microphone sensors located on the robot. RiSH was able to distinguish human activity sounds (e.g., eating, drinking, brushing teeth, running water, snoring, talking) from falling sounds for scripted tasks performed by 12 subjects. As their experiments demonstrated, the ability to detect falls

improved when both audio and wearable data were utilized. Fall detection reached 80% for the testbed experiments.

While researchers foresee a need for robotic activity assistance in daily environments, many wonder whether older adults and individuals with health needs would be accepting of such assistance, particularly from an automated robot. Perhaps surprisingly, several recent studies indicate that these individuals are interested in the benefits these systems can provide [32]. In a study led by Begum et al. utilizing robots to assist subjects with hand washing and making tea [33], older adults shared that they do want the help that can be provided through robotic assistance, yet they do not want a robot in their home. In this study, experimenters monitored subjects as they perform two scripted activities, wash hands and make tea. A tele-operated robot assisted subjects through navigational guidance and conversational interactions. The robot led subjects to the kitchen when it was time to make tea, answered task-specific questions, and led subjects back to their caregiver when the task was complete. In contrast with the older adults subjects, caregivers more uniformly expressed a desire to install robots in the homes of loved ones to provide assistance with these tasks.

As Mitzner et al. point out [34], robots will be more widely accepted if they are more carefully designed to meet the needs of older adults. Because robots are becoming more reliable, researchers have had opportunity to obtain feedback on actual human-robot interactions with this target group and have provided guidance on impactful design. For example, Hoffman et al. [35] caution that robot movements and gestures are socially expressive, thus these movements affect the level and type of engagement humans will have with them. Destephe et al. [36] as well as Bisio et al. [37] note that humans better understand robot intent if the robot moves and gestures in ways that actually mimic biological movement patterns. As a specific example, Kupferberg et al. [38] found that robot movement velocity should be similar to that of nearby humans, because these similar speeds facilitate perception of a humanoid robot as a true interactive collaborator in their daily needs. Some efforts have focused specifically on being sensitive to a human's interactions needs. For example, the robot Pepper [39] perceives its owner's emotions and adapts its behavior accordingly. As Pepper looks at the human, waves its hand, or shakes hands with a human, its own head angle, speed of movement, and reaction to human gestures is guided by reinforcement learning to be pleasing to the humans around it.

Another decision that affects CPS design is how much assistance to provide to individuals with health needs. Honda [40] has designed their humanoid robots to provide high levels of assistance automation, including fetching food and controlling household devices. Mataric et al. [41], however, realize that in many cases robots should be designed to help people be independent and thus assist them only when needed, not replacing the capabilities of humans to perform tasks when they are able.

The type of activity assistance that robot cyber-physical systems provide to individuals spans a broad range. In some approaches, robots play a passive role, monitoring individuals as they go about their daily routines. The robot designed by Goher et al. [42] monitors activities to provide medicine reminders and update caregivers of related anomalies.

In other work, robot systems provide more physical, hands-on assistance to individuals with physical limitations. Liao et al. [43] created a robot that provides physical support for upper-arm motions that are required in daily tasks, helping those with upper-arm weakness to complete such tasks independently. The robot from Riken [44] also provides physical support by transferring individuals between their bed and wheelchair when the task cannot be accomplished on their own. Fasola and Mataric [45] designed a robot that acts as an exercise coach, providing physical assistance as well as feedback to individuals needing physical therapy. When assistance requires more social engagement than physical support, robots from Pu et al. [46] and from Wada et al. [47]

can be employed. These robots elicit feelings of joy and relaxation from their owners, helping in critical moments when individuals are lonely or agitated.

The work we introduce in the RAS system builds on elements of these previous projects. RAS does utilize sensing, activity monitoring, and robot-human interaction to aid individuals in performing daily activities. The greatest similarity is with some recent efforts that utilize robot cyber-physical systems for activity guidance in fairly constrained settings. As an example, Bovbel and Nijat [48] designed a robot cyber-physical system that aided individuals with cooking tasks in home settings. When activated, the robot finds the resident, leads them to the kitchen, then points to items stored in pre-programmed locations that they will need to complete a selected recipe. Hidden Markov models capture past human movement patterns, thus reducing the effort needed for the robot to locate and approach the resident. The number of rooms that were searched was smaller on average than using random search.

RAS represents a cyber-physical system that integrates many components. Some of the components, such as activity identification and error detection, represent novel enhancements to the field. The robot hardware is commercially available, and software for object detection and robot localization were based on research from other groups that we enhanced to improve performance of our tasks. The main contribution of this work, however, is the melding of these individual technologies to create a fully automated smart home/robot activity assistant. By automating activity recognition and tracking, robot navigation, object detection, and human-robot interface, we demonstrate that RAS is able to provide support in actual homes with residents performing activities that occur in their normal daily routines. Because RAS is fully implemented and tested in home settings for routine activities, we also obtain feedback on the usability of the cyber-physical system that can inform future research and commercial efforts.

3 METHODS

To provide assistance with everyday activities, RAS needs to sense and identify activities as well as detect and respond to activity errors. Here, we discuss each of the components. These include the smart home, which provides sensing capability, activity recognition that identifies the activities in real time, and software that detects activity errors. We then describe the capabilities of the RAS robot, which offers activity assistance in the form of prompting, video reminders, and object retrieval.

3.1 Sense: CASAS Smart Home

In our activity support system, the first cyber-physical element is a smart home. In our CASAS smart home, ambient sensors are embedded into an existing residence. Data are collected continuously from the sensors while residents go about their normal routines. As shown in Figure 2, sensors include passive infrared (PIR) motion sensors (coupled with ambient light sensors), magnetic door sensors (coupled with ambient temperature sensors), accelerometer-based Estimote item sensors, and pressure-based items sensors. All sensors send readings via Zigbee when there is a change in state (i.e., when there is new movement or movement has stopped, when a door opens or closes, when an item is accessed, or when the light or temperature levels change more than a threshold value). Door sensors are placed on regularly used external doors and cabinets holding key items such as medicine dispensers. Downward-facing motion sensors are placed in functional spots of the house with a one-meter field of view, and other motion sensors are angled to monitor an entire room or large area.

Ambient sensors placed in a smart home are discrete event sensors. As such, they do not provide readings at constant sample rates but instead send readings with the sensed state changes. For example, a motion sensor sends an “ON” message when it senses new motion in its field of view

and sends an “OFF” message when motion has ceased. The number of readings therefore increases when the activity level rises and decreases when less activity occurs (e.g., when the residents are sleeping). Messages sent from ambient sensors are sent to a Raspberry Pi, which attaches timestamps and sensor identifiers to each message before storing it in a secure SQL database. While amount of collected data varies depending on the number or installed sensors and amount of activity that occurs in the home, an average of approximately 3K sensor events are collected per day in the homes. Although additional sensor sources can be integrated into the system such as wearable and video data, currently RAS utilizes ambient sensors for its primary data source.

3.2 Identify: Activity Recognition

The second element of the RAS system is real-time recognition of activities from ambient sensor data. This element distinguishes this work from earlier research on assisting individuals with targeted activities in scripted settings such as hand washing, cooking, and making tea [49–51]. In our work, RAS recognizes the activity that an individual performs in real time as it is performed. Based on the detected activity, RAS retrieves the corresponding information about how the activity is performed, tracks activity steps, and calls the robot to action if an omitted step is detected. In this way, assistance can be provided for any recognizable activity in a person’s home environment.

Activity recognition can be characterized as a supervised machine learning problem. Here, a machine learning algorithm is given a set of (x, y) pairs where x represents a data point and y represents a corresponding class label. In the case of activity recognition, data is collected in the form of a series of n sensor events (or sensor readings) $s = \{e_1 e_2 .. e_n\}$ from which descriptive features $x = f(s)$ are extracted. The machine learning algorithm maps x onto a value from a set of A possible activity labels, $y \in \langle a_1, .., a_A \rangle$. Activity recognition has been widely researched for multiple sources of data, including ambient sensors [52], wearable sensors [53], and other information sources such as video [54]. If data is pre-segmented into non-overlapping activities, then each segment can be mapped onto an activity label. Because the smart home data is not pre-segmented

19:20:48.98	LoungeChair	On	19:21:09.72	Bedroom	On	19:21:22.86	Entry	On
19:20:50.99	LoungeChair	Off	19:21:14.44	Bedroom	Off	19:21:24.28	OutsideDoor	Open
19:21:01.11	Bedroom	On	19:21:15.55	LoungeChair	On	19:21:25.69	Entry	Off
19:21:02.98	Bedroom	Off	19:21:17.62	LoungeChair	Off	19:21:28.14	Entry	On
19:21:03.55	Bedroom	On	19:21:18.18	Entry	On	19:21:33.12	OutsideDoor	Close
19:21:04.71	Bedroom	Off	19:21:21.93	Entry	Off	19:21:34.32	Entry	Off

Fig. 3. Sample of data sequence containing 18 events. Each event is described by the time of day, the sensor identifier, and the sensor message. Dates have been removed here to preserve anonymity. AR labeled this sequence as “Enter home.”

and we need to recognize activities in real time, we instead map a sequence containing the most recent sensor events onto an activity label. For our experiments, the length of each sequence is 30 sensor events. The output activity label can be interpreted as the activity that is performed at the end of the fixed-length sequence. This way, activity labels are generated that indicate the most recent activity being performed. Figure 3 provides a sample of collected data and Table 1 summarizes features that are extracted from these sensor event sequences.

To train the learning algorithm, ground truth activity labels must be provided for some of the data. External annotators provide this information based on a floor plan, sensor layout description, input from the residents on when and where they typically perform activities, and the raw sensor data. Inter-annotator agreement (a measure of labeling consistency) is $\kappa = 0.80$ for the 12 activities that we model: bathe, bed-toilet transition, cook, eat, enter home, leave home, personal hygiene, relax, sleep, take medicine, wash dishes, and work. One challenge that has been issued for health CPS is to perform activity recognition with minimal training for each new home and person [55]. Toward this goal, we convert each unique sensor identifier to a functional area description (see Figure 3) and learn a single population-generalizable model based on one month of annotated sensor data from each home. For the experiments in this article, the model is based on data collected from 55 single-resident smart homes. This dataset contains a total of 19,347,964 sensor readings. Recognition of the 12 activities for this data is 98% across all of the collected data using three-fold cross-validation. While the experiments are based on single-resident settings, the approach can also be used for activity recognition in multiple-resident settings. In these settings, sensor events are still labeled with corresponding activity categories, although currently the technology does not track the identity of the residents that are performing activities.

3.3 Assess: Activity Error Detection

While activity recognition is becoming an established field, less attention has been devoted to tracking activity steps to detect errors during the performance of an activity. In earlier work, we explored detecting errors as deviations from the normal strategy for an activity [56]. Previous literature characterizes activity errors that are frequently committed by individuals with cognitive impairment. An *omission* error occurs when a step necessary for accurate task completion is not performed [57]. A *substitution* error occurs when an abnormal object or location is used for the activity and thus disrupts activity completion, as might occur when coffee grains are mistakenly used to make tea [58]. An *irrelevant action* is one that is performed during an activity but is unrelated and unnecessary for the activity, while an *inefficient action* is one that slows down or compromises the efficiency of activity completion. Additional activity difficulties include *perseveration*, in which an individual engages in an activity even after it is completed such as adding extra unneeded ingredients while cooking; and *searching*, in which an individual wanders the home or repeatedly opens and shuts cabinet doors in search of an activity-related item.

Table 1. Features Extracted from Ambient Sensor Data for Activity Recognition

Raw Data	Type	Description
	MA1..MAa; M1..Mb	Passive infrared motion sensors (broad area or focused region)
	D1..Dc	Magnetic door sensors
	LS1..LSd	Ambient light sensors
	T1..Te	Ambient temperature sensors
	Item1..Itemf	Accelerometer-based item sensors
	Date	month, day, year
	Time	hour:minute:second.ms
Features	Time of the most recent sensor event	Hours past midnight; seconds past midnight; day of week
	Sequence size	Duration in seconds
	Elapsed time since most recent sensor event	Seconds
	Dominant sensor	Sensor identifier for current and previous sequences
	Identifier for most recent event	Integer
	Location for most recent event	Integer
	Location for most recent motion sensor event	Integer
	Sequence complexity	Entropy based on sensor counts
	Change in activity level between first half and second half of sequence	Continuous
	Transitions between locations within sequence	Integer
	Number of distinct sensors in sequence	Integer
	Event counts for each sequence	(a+b+c+d+e+f) integers
	Elapsed time since last event for each sensor	(a+b+c+d+e+f) seconds

A novel contribution of this work is automated detection of these types of errors that have previously required experimenter monitoring. In the current RAS system, only omission error detection is implemented. Previous research found that omission errors uniquely predict a person's ability to be functionally independent despite cognitive and memory limitations [57]. In this previous study, examples of omission errors included failure to retrieve a broom for a sweeping and dusting task (a critical omission) or failure to turn off the television at the end of watching a DVD (a non-critical omission). Additionally, RAS currently assumes that activities are not interrupted—one activity is performed to completion before another starts. In RAS, a directed acyclic graph (DAG) is constructed for each activity, indicating the order in which activity steps should be performed based on observation of normal behavior. During activity monitoring, RAS identifies the current activity then traverses the corresponding DAG as the activity steps are tracked. The activity can be identified as described in the previous section even when the first step is skipped, which is interpreted as an indication that the individual wants to engage in that specific activity. When steps are detected in a non-DAG-supported order, RAS signals that an activity error has occurred. This signal may be used to initiate an appropriate intervention.

Fig. 4. RAS software, hardware, and communication components. Components are located within a CASAS smart home, a CASAS server, the RAS on-board computer (Joule or Jetson), robot hardware, or an Android tablet interface. Components communicate using RabbitMQ or ROS.

3.4 Act: Robot Intervention

In the RAS system, a robot intervenes by approaching a smart home resident when an activity error is detected, offering to help the resident by showing a video of the entire activity, of the missing activity step, or leading the resident to an object that is needed for the step. Performing these tasks requires that the robot be mobile, be able to navigate about the space, detect and store locations of human residents and key activity-related objects, and have an interface that is easy for individuals with memory impairment to use.

RAS contains many hardware, communication, and software components that must work together as a whole. These components and their connections are highlighted in Figure 4. RAS tasks are divided among the CASAS smart home (sensors and servers), the on-board robot computer, the Nvidia computer, the tablet app, and the robot hardware. Communication between these components is driven by either the RabbitMQ message-broker software (if communicating with CASAS) or ROS (if communicating internally within robot software elements). While the RAS software configuration shown in Figure 4 is unique to this project, many of the components are commercially available, such as the robot hardware platform, RabbitMQ, ROS, object detection software, and SLAM software.

Robot task management is specified by two task state machines, corresponding to the tasks “go to person” and “go to object.” These two state machines are designed specifically for RAS but are implemented using the existing ROS State Machine infrastructure, or SMACH. Within each SMACH, a node represents a single function and transitions are made based on the outcome of the function. Figure 5 shows the two SMACH state machines we constructed for RAS. When the “go to person” SMACH is initiated, RAS starts in the “FindPerson” execution state. Within this state, RAS queries the database for time and location where the human resident was last located. If the timestamp is at least 10 seconds old, then RAS scans the space for the robot by rotating the

Fig. 5. RAS SMACH state machines for the two robot tasks: Go to Person (left) and Go to Object (right).

camera. Once the person is located, the corresponding location is updated in the database and the SMACH transitions to “GotoLoc,” which brings up RAS’s navigation component. If the human is not located after scanning the entire space, then the RAS robot moves to a new position and repeats the search. RAS employs a similar SMACH to find and navigate to desired objects. The current RAS implementation pre-specifies the objects to detect and store in a location database; future expansion of the research may automatically detect important objects based on past execution of the activity and thus more greatly adapt to each person’s method of performing daily activities.

3.4.1 Robot Hardware. The RAS robot (see Figure 6) is built on a TurtleBot 3 base, which includes 360° LiDAR for simultaneous localization and mapping (SLAM) as well as navigation, an on-board Intel Joule computer, wireless communication, and the open-source control module (OpenCR) for the robot operating system (ROS) to process sensor data and control the wheels. We made a number of unique modifications to the base robot for RAS operation. To improve mobility, stability, and payload, we 3D printed “omni wheels” for the back of the robot. These wheels allow the robot to move in all directions. Human and object detection is performed by an Orbbec Astra RGBD camera heightened by a 4 ft mast of lightweight aluminum while two servos atop the mast facilitate pan-tilt control. Captured video is processed by an Nvidia Jetson computer, which also maintains a central map of the smart home as well as the locations of objects and humans within the space. To provide an interface for smart home residents, we attached an Android tablet with our RAS-interaction app to the mast.

3.4.2 Robot Navigation. The two SMACH state machines contain execution states that require navigation of the mobile robot to a specified location in the smart home. As mentioned earlier, the RAS system needs to operate in any smart home environment. This means that RAS must map the space as well as localize the robot and chart navigation paths. Cartographer [59] generates a sensor-based map of the smart home layout (see Figure 7) and continuously updates the robot’s position on the charted map. One setup step is required: A human operator initially drives the robot around the space to ensure that the map includes all navigable smart home regions.

As Figure 7 illustrates, large pieces of furniture are mapped by Cartographer. During navigation, the robot’s LiDAR sends additional information about obstacles. These two sources of information are combined into a cost map indicating locations that will likely cause robot collisions. A “safe zone” buffer is built around each detected obstacle and is avoided during navigation, creating paths that achieve maximum clearance on all sides of the robot.

3.4.3 Human/Object Detection. There are two situations in which RAS performs object detection. First, when an activity error is detected, the robot needs to find, approach, and face the human (human detection). Second, if a person forgets where an object is located that is needed for an

Fig. 6. The RAS robot.

activity (e.g., a medicine dispenser), RAS leads the person to the object. This feat requires object detection as well as the ability to store object locations in the map and navigate to the intended location. Video data for object detection is collected with an RGBD camera that is poised at the top of the robot's mast (see Figure 6). The RGB channels obtain color distributions and the D channel collects depth, or distance to the nearest object from each pixel in the image. In addition to assisting with better visual coverage of the space, the Arduino-controlled mount reports current pan-tilt angles to RAS. These servo angles facilitate calculation of the location where the camera is pointing relative to the robot. By varying the mount angles, RAS can scan the entire space. Using a coordinate transformation, RAS can express the object's location in terms of map coordinates to store for later use.

The goal of object detection is to generate a bounding box around each object of a given class (e.g., each medicine container) within a captured image. As illustrated in Figure 8, once the bounding box is created, 3D positions of each bounding box are calculated based on the calibrated depth channel. Next, the 3D position relative to the camera is converted into a position relative to the navigation module's map. This conversion is performed by applying a coordinate transform to the original 3D location. Finally, we update the last-seen location of the object and save it to a database that RAS can later query when the object is needed.

We employ a convolutional neural network (CNN) that has demonstrated success for general object detection [60]. Here, training points are fed to the network consisting of bounding box-constrained regions of a specific image, where different bounding box specifications define

Fig. 7. Actual layout for bottom floor of smart home with added furniture from Figure 2 (left) and RAS-learned map of navigable smart home regions (right). In the robot map, green dots highlight the current LiDAR scan and the dark green rectangle represents the robot's bounding box.

different data points. The network predicts a class (object label) for each box as well as an amount and direction to shift the box. During training, we use an objective function that balances minimizing error for incorrect class labels and minimizing error for incorrect bounding box positions [61]. Utilizing a complex, multi-objective function introduces numerous hyperparameters that need to be selected to define the network structure, number of bounding boxes, and bounding box constraints for size and position.

We train a specific type of CNN, a region-based fully convolutional network (R-FCN), using a combination of images of RAS-selected objects collected from the smart home and additional images from online databases. The R-FCN separates object detection into two steps. In the first step, a portion of a feature extractor network proposes boxes based on the input image that may contain the objects. After cropping the image to contain only the boxes, the rest of the network processes the updated features in the second step to classify the objects. Each image is preprocessed by manually drawing a bounding box around the object of interest and assigning a class label. Once trained, RAS can perform object detection in real time using video from the robot camera. RAS processes approximately 10 frames per second using this approach, which supports looking for objects of interest at the normal speed of the robot. The camera's depth sensor generates a point cloud that indicates the 3D position of a detected object relative to the camera. After transforming the relative position to an object location position, the object's location is updated in RAS's object location database. The same process is applied to maintain most-recent locations for smart home residents as well.

3.4.4 Robot-human Interface. The RAS robot-human interface was created with guidance from an open discussion with the design team and a group of 25 potential participants. Interactions

Fig. 8. RAS human/object detection. Images that may contain objects of interest (e.g., flower pot, medicine dispenser, glass, dog) are collected from robot camera images or publicly available image repositories (upper left). These training images are manually preprocessed to label objects and indicate the corresponding object bounding boxes (upper right). Once labeled, the training images are fed to a convolutional neural network (lower left). The trained CNN is stored in RAS to use for real-time object detection based on video collected from the RGB camera (lower right).

are facilitated by an Android app that maintains three expressive modes. The first mode, **monitor mode**, is adopted when RAS is passively tracking activities within the house (Figure 9, top). If the smart home detects an activity error, then RAS transitions to **assistant mode** (Figure 9, bottom). In assistant mode, RAS initiates the Go to Person SMACH and approaches the smart home resident with an offer of help. If RAS's offer of help is accepted, the tablet then offers four options (Figure 10): show a video of the entire activity on the tablet, show a video of the current (missing) activity step, lead the resident to the object that is needed for the activity step (the name of the object is completed by the interface), or confirm that the step has been completed. Once the error is corrected, RAS transitions to **success mode**, offering a pleased expression to the user (Figure 6).

All of the activity videos are prerecorded and segmented, ideally showing the resident performing the activity in their own home. Videos are stored on the CASAS server and played through a web server interface when needed. If the resident requests that the robot lead the way to an object needed for the activity, then the tablet plays a "Follow me" audio clip and RAS initiates the Go to Object SMACH. Once the object is in visual range and physical reach, the robot stops and the interface plays a "Here you go" audio clip.

Fig. 9. The tablet shows a passive avatar expression when RAS is in monitor mode (top). In assistant mode, the tablet shows an inquisitive expression while asking the resident if they want help with the activity (bottom).

4 EXPERIMENTAL PROTOCOLS AND RESULTS

We evaluate the RAS cyber physical system to assess our three goals: ability to track activities across different homes, different residents, and even different types of sensors; usability of the system by these different individuals; and capability of performing in natural home situations. To this end, we evaluate RAS in two types of settings. First, we conducted a controlled experiment in a single smart home with younger adult and older adult participants and item sensors. Second, we deploy RAS for multi-day use in two smart homes to support routine activities.

Fig. 10. The tablet interface offers four responses: “Show video full” (of the entire activity), “Show video step” (of the current activity step), “Take me to the object,” and “I did it!” (top). If an activity video is requested, then it is played on the tablet interface (bottom).

4.1 Testbed Experiment

Our first experiment setting occurred in an on-campus smart apartment. We used the first floor of the smart home with the sensor layout that is shown in Figure 2. The furniture layout and corresponding RAS navigational map is shown in Figure 7. We recruited $n = 54$ participants (27 younger adults and 27 older adults) to visit the smart apartment, one at a time, and perform a set of scripted activities. Each activity was reflective of normal routine behavior: prepare to take the dog for a walk, take medicine with food, fill the watering can with water and water

Fig. 11. Estimote sensors attached to house keys and a flower pot (left). Passive infrared motion sensors installed on the ceiling of a smart home (right).

plants in two locations around the home. Additionally, these activities represent instrumental activities of daily living (IADLs) that are commonly assessed by IADL questionnaires as well as by performance-based measures of everyday competency [62]. Successful completion of IADLs requires intact cognitive abilities such as memory and executive functions [63]. Therefore, these activities are likely candidates for needing assistance in home settings.

For each activity, subjects performed the activity four times, once without errors and three additional times with missing activity steps. In the case of walking the dog, possible errors were failing to retrieve an umbrella, failing to grab house keys, and not bringing the dog leash. In the case of taking medicine, errors include forgetting to take food with the medicine, not retrieving the medicine, and not taking the medicine. For the water plants activity, errors included not filling the water can and forgetting to water some of the plants. Steps were chosen for omission based on a corresponding interaction with an object. This is a feature that distinguishes instrumental activities of daily living from basic activities of daily living. Researchers have shown that declining ability to perform IADLs is related to early symptoms decline in cognitive abilities [64], so we want to focus on this distinction for the RAS use cases. Some participants were unable to complete all activity tasks in the available time and others performed activities multiple times. Data were collected for a total of 683 activity occurrences.

In this experiment, we use Estimotes [65] to collect activity data. Estimotes are stickers with embedded accelerometers and temperature sensors (see Figure 11). We attached them to objects throughout the smart apartment. Because the sensors provide raw acceleration data, we determined that an object was being moved or manipulated when the acceleration values exceed a specified threshold. This information is sent to the CASAS, which timestamps the object usage information and stores it in the CASAS database.

Table 2 summarizes the performance of RAS for this experiment. For each entry in the table, malfunction rate is computed as the ratio of activity occurrences that include the particular type of malfunction to the total number of activity occurrences. As Figure 4 shows, there are many components that need to operate smoothly for RAS to function. As a result, we summarize system error rate based on the corresponding categories of errors that may occur, as follows:

- False positive activity error detection (FP). These are errors that were detected by RAS but did not actually occur.
- False negative activity error detection (FN). These are errors that were committed by subjects but were not detected by RAS.
- Object detection malfunction. If a smart home resident requests that RAS lead the way to an important object for an activity, then RAS initiates the Go to Object SMACH. The robot

Table 2. Malfunction Rates, Categorized by Type of Malfunction

Error type	Total
FP	0.008
FN	0.005
Object detection	0.000
Human detection	0.088
Navigation	0.054
Interface	0.021
System	0.029
Experimenter intervention required	0.455

navigates to the stored object location if one is available, then scans the space for the object using object detection.

- Human detection malfunction. Once an activity error is detected, RAS signals the robot to approach the smart home resident via the Go to Person SMACH. A malfunction is reported when the robot does not move within an acceptable (two-foot) range or does not turn to directly face the human.
- Navigation malfunction. The Go to Person and Go to Object SMACHs rely on the navigation module to successfully steer the robot to the desired location without incident. However, navigation malfunctions did occasionally occur. These included the robot getting stuck on an object or halting movement prior to reaching the intended location.
- Interface malfunction. App interface malfunctions were noted on a few occasions by the experimenters. These include the tablet not registering the user response or not providing a verbal prompt to follow the robot.
- System malfunction. These malfunctions reflect broken communication between the CASAS server and RAS manager or between the RAS manager and the robot.

Activity errors were detected in this experiment with a sensitivity of 0.955 and a specificity of 0.992. However, other types of RAS assistance problems did occur. While the malfunction rate for any given category of system issue was low (as shown in Table 2), there were still situations that required experimenter intervention. In some cases, intervention was required when a RAS component malfunctioned. In other cases, intervention was required because a subject went off-script with their routine and introduced anomalies that were not handled by RAS. The independent-success rate for RAS was 0.600. This number indicates that even though individual RAS components may be fairly robust, the overall system requires additional refinement to be sensitive to the dynamic nature of human activities in their own home environments. In this scripted experiment, the smart home resident would independently complete activities with a success rate of 0.250. With RAS intervention, the expected success rate improves to 0.700.

To further assess the usability of RAS, we asked the younger adult and older adult participants to provide feedback on the system usability and usefulness after finishing the scripted activities. Specifically, they completed a Post-Study System Usability Questionnaire (PSSUQ). PSSUQ uses Likert ratings, where 1 = strongly agree and 7 = strongly disagree (lower numbers are more favorable). Table 3 summarizes the results from the two groups of participants.

4.2 In-home Experiment

We next evaluate the ability of RAS to operate for multiple days, without experimenter intervention, in a smart home assisting a resident with daily activities. This experiment differed from the

Table 3. PSSUQ Mean and Standard Deviation Results, by Group and Overall

Questionnaire	Younger adults (n = 27)	Older adults (n = 17)
PSSUQ		
Overall	4.55 (1.95)	4.38 (1.92)
System usefulness	4.37 (1.89)	4.22 (2.07)
Interface quality	4.83 (1.94)	5.06 (1.70)
Information quality	4.66 (2.23)	5.50 (1.72)

previous one in several ways. First, RAS operated autonomously with no experimenter assistance. Second, subjects selected activities for assistance that were part of their normal daily routine and were performed in their routine manner. Third, activities were recognized and tracked using ambient sensors (passive infrared and door) instead of Estimote item sensors (the two types of sensors are illustrated in Figure 11).

After receiving feedback for the testbed experiment, we modified RAS to reduce communication delays with the CASAS server. Additionally, we modified robot navigation to keep the robot farther away from detected objects, reducing the likelihood of collision with furniture. Finally, we performed activity tracking on the CASAS server rather than on the local robot computer, further reducing network traffic. This modified version of the robot was employed for the in-home study.

We performed this experiment in two different smart homes. Home 1 housed a younger adult female resident and Home 2 housed an older adult couple (only one resident interacted with the robot and provided feedback). In each home, we installed ambient sensors and collected data for one week to train activity recognition models. During this initial visit, experimenters met with residents to select three activities they performed multiple times each day for which RAS could provide support. Experimenters explained to residents the purpose of RAS and its goal to assist with basic and instrumental ADLs. The resident then suggested activities they performed routinely that fit in these categories. Experimenters recorded video of the resident performing each activity to later use for prompting videos and the RAS robot learned a map of the space.

After activity models were learned, the RAS robot was installed in the smart home to provide activity support for 3–4 days. During this time, participants performed the selected activities at least twice each day while RAS monitored activity performance. Participants were instructed to inject one activity omission error for each activity one time each day, allowing RAS to aid them in finishing the incomplete activity. In each case, the participant chose the step to omit. To motivate their decision they considered a type of omission they would be likely to commit and want assistance to complete. Rather than skip different steps during different performance of the activity, the same step was skipped each time. The activities were not scripted and no experimenter was present to guide the resident. As a result, the activities naturally varied from one day to the next. RAS performance was averaged over these multiple variations of the same activity and same error to assess the generalized performance of the system. The selected activities with corresponding steps and errors are summarized for each home in Table 4. Participants were given log books to record notes about interactions with RAS. After the multi-day RAS support, participants completed system satisfaction questionnaires similar to those completed in the testbed experiment.

For these in-home tests, participants maintained a log book of their RAS interactions. In addition to providing notes on the daily activities and RAS support, participants provided quantitative feedback after each activity, rating the RAS interaction. Once the study was completed, participants answered an additional set of survey questions indicating their overall rating of the system.

Table 4. Activities Selected for RAS Support in Two Smart Homes

Home 1	Home 2
Activity: Eat	Activity: Take medicine
Bring food and dishes to dining room Eat meal <i>Return items to kitchen and wash dishes</i>	Retrieve pill crushing tools, crush medicine Fetch applesauce from refrigerator, mix with pills Take medicine Wash tools in kitchen sink <i>Return tools</i> Put remaining applesauce in refrigerator
Activity: Take medicine	Activity: Make coffee
Retrieve medicine <i>Fill glass with water</i> Take medicine with water Return medicine Return glass	Retrieve coffee maker Pour water into coffee maker Retrieve coffee grounds, scoop into coffee maker <i>Put coffee grounds back</i> Put coffee maker back
Activity: Work	Activity: Eat
<i>Retrieve computer</i> Sit on couch to work on computer	Retrieve dishes Assemble food Take items to dining room table and eat <i>Wash dishes in sink</i>

Steps that are omitted during activity error are highlighted in green italic font.

Table 5. Survey Responses for In-home Study. Feedback uses scale 1 (extremely dissatisfied) – 7 (extremely satisfied)

Question / Scale	Home 1	Home 2
Ease of completing activity (daily, averaged)	6.83	4.38
Mistake caught in time to fix (daily, averaged)	4.00	6.00
Robot offers enough support for day-to-day activities (daily, averaged)	6.13	4.67
Satisfied with robot and its help (daily, averaged)	5.58	3.50
I was able to complete the activities using RAS (overall)	6	6
I felt comfortable using this system (overall)	6	7

A summary of participant responses is provided in Table 5. As the results indicate, RAS was able to guide residents through activities when errors were detected. However, participant notes indicate that the current system is inflexible and needs to handle a greater variety of activity contexts to be broadly useful.

5 DISCUSSION

Evaluation of RAS in a testbed apartment as well as two homes with participant residents reveals both successes and challenges of deploying a robot/smart home activity support system. Because the RAS system has many components, there are many sources of possible error. For the testbed experiment, the false positive and false negative activity error detection malfunction rates indicate that in 9 of the 683 activity occurrences, activity errors were incorrectly detected. Experimenters noted that these false positives and false negatives were attributed to the Estimates. The Estimate sensors were attached to items throughout the apartment. If an item was moved, then the Estimate

software sent a “Move” message that was collected by RAS. Because the amount of movement needed to exceed a specified threshold, in some cases the sensors were too sensitive and in other cases they were not sensitive enough.

A second source of error is object or human detection. In a separate study, we evaluated our software for its ability to detect the specific objects used in the study as well as to detect humans under a variety of lighting conditions. The evaluation indicated a macro-average object detection precision of 0.99, which is consistent with the error rates we observed in the testbed experiment. However, human detection errors were more common because of the much greater variation in human images. Training the detector with a much greater number of human faces, varied lighting conditions, and different background settings may improve the results.

During the testbed experiments, RAS did experience 37 navigation errors. In 10 of these cases, the robot got stuck on an object. Typically this occurred at the corner of the object and could be addressed by adjusting the SLAM software to keep the robot farther away from all objects in the space. Because the goal was to move robots to a distance that is considered commonly accept for one-on-one interactions, any situation in which the robot halted greater than two feet away from the human was considered an error. In the future, we may modify this measurement to allow the participant to indicate when the interaction distance is comfortable.

The remaining malfunctions were highly varied. In some cases, participants did not push a correct button on the interface, did not push hard enough, or accidentally reset the app, causing an interface malfunction. In other cases, the tremendous amount of wireless traffic near the testbed site caused message interference. The smart home testbed is part of a large apartment building with numerous wireless networks. The amount of traffic surrounding the apartment, as well as within the smart home, caused some network packets to be dropped. Future versions of the RAS system will need to be made more robust to handle these situations more gracefully. Because the popularity of smart homes is growing, robust communication will become increasingly important to ensure reliability of all related services.

For the in-home study, RAS committed one false positive error detection (Work) and one false negative error detection (Eat) in Home 1. These errors were likely due to the coarse-granularity detection of resident movement and activity in the confined space where all of the activities were performed. Home 2 presented a greater activity tracking challenge, because multiple residents live in the home. While activity recognition still labels each sensor event with a corresponding activity label regardless of the resident who performs the activity, more movement throughout the house increases the complexity of the recognition task. As a result, RAS committed three false positive error detections, all for the Eat activity.

Interestingly, RAS also correctly detected an error in each house that was not part of the scripted tasks or preplanned errors. In Home 1, RAS twice reminded the resident to clean up after a meal when the resident was slow to do this. In Home 2, RAS prompted the resident to return medicine tools and coffee to their storage locations throughout the day when the resident did not do these steps. Combined across the in-home experiments, activity errors were detected with a sensitivity of 0.905 and a specificity of 0.988. RAS activity assistance success (defined as the ratio of activity occurrences with successful error detection, intervention, and navigation to the total number of activity occurrences) was 0.84 in Home 1 and 0.83 in Home 2. In these homes, given the human-only success rate of 0.500, RAS assistance would improve successful activity completion to an expected rate of 0.915.

For both the testbed experiment and the in-home experiments, participants found interacting with the robot generally enjoyable and the interface easy to understand. However, participants frequently felt that the robot moved too slow. However, participants occasionally would move out of the path of the robot if the robot approached them too quickly. In these situations the participants

viewed the move as too aggressive and were not certain of the robot's intentions. These observations indicate that additional research is needed to ensure that all of the robot movements are natural, non-threatening, yet efficient. These interactions will become more complex as the number of people in the home increases and when the residents entertain visitors who are unfamiliar with the robot movements and interactions.

6 CONCLUSIONS

In this article, we introduce RAS, a cyber-physical approach to offering activity support for individuals with health limitations. By partnering the physical capabilities of smart home sensors and robot movements with the computational elements needed to track activities and detect objects, RAS aims to provide complete activity support. We validated in the article that the system can sense human behavior, identify current activities, assess whether an activity error has occurred, and intervene by playing a video of the activity or leading the resident to a needed item. Through a testbed and two in-home experiments, we demonstrated that the numerous sensing, computing, and actuating components can work together to provide activity support that was accepted by younger and older participants.

Future work on this cyber-physical system will focus on improving communication robustness and increasing the diversity of activities to support. RAS can also mature by handling interwoven activities and additional types of activity errors. Additionally, we would like to enhance the system to include automated video recording of activities as well as segmentation into task steps. We anticipate that the RAS robot can automatically capture video of the resident performing activities to use as future intervention cues. By incorporating change point detection techniques, video and smart home data can be partitioned into individual activity steps for monitoring and prompting. This will reduce the manual overhead of using the system and allow the technology to be used in a greater variety of naturalistic settings for individuals who need activity support to remain functionally independent.

ACKNOWLEDGMENTS

The authors would like to thank Shivam Goel, Sepehr Nesaei, Gabriel de la Cruz, Matthew Taylor, Asim Fauzi, and Julia Maliauka.

REFERENCES

- [1] Future Market Insights. 2018. Cyber-physical system market: Global industry analysis (2013-2017) and opportunity assessment (2018-2028). <https://www.futuremarketinsights.com/reports/cyber-physical-systems-market>.
- [2] Web of Science. 2019. Web of Science: Search=cyber-physical systems; sort=year. 2019.
- [3] J. Iriando and J. Jordan. 2018. Older people projected to outnumber children for first time in U. S. history, 2018. <https://www.census.gov/newsroom/press-releases/2018/cb18-41-population-projections.html>.
- [4] P. Buonocunto, A. Giantomassi, M. Marinoni, D. Calvaresi, and G. Buttazzo. 2018. A limb tracking platform for tele-rehabilitation. *ACM Trans. Cyber Phys. Syst.* 2, 4 (2018), 30 2018.
- [5] J. A. Ramirez-Bautista and J. A. Huerta-Ruelas. 2017. A review in detection and monitoring gait disorders using in-shoe plantar measurement systems. *IEEE Rev. Biomed. Eng.* 10, 299-309, 2017.
- [6] M. Mancini et al. 2016. Continuous monitoring of turning mobility and its association to falls and cognitive function: A pilot study. *J. Gerontol. Ser. A Biol. Sci. Med. Sci.* 71, 8 (2016), 1102-1108, 2016.
- [7] T. L. Hayes, T. Riley, N. Mattek, M. Pavel, and J. A. Kaye. 2014. Sleep habits in mild cognitive impairment. *Alzheimer Dis. Assoc. Disord.* 28 (2014), 145-150.
- [8] G. Sprint, D. Cook, R. Fritz, and M. Schmitter-Edgecombe. 2016. Detecting health and behavior change by analyzing smart home sensor data. In *Proceedings of the IEEE International Conference on Smart Computing (SMARTCOMP'16)*.
- [9] S. Difrancesco et al. 2016. Out-of-home activity recognition from GPS data in schizophrenic patients. In *Proceedings of the IEEE International Symposium on Computer-based Medical Systems*. 324-328.

- [10] M. Boukhechba, Y. Huang, P. Chow, K. Fua, B. A. Teachman, and L. E. Barnes. 2017. Monitoring social anxiety from mobility and communication patterns. *Proceedings of the ACM International Joint Conference on Pervasive Ubiquitous Computing*. 749–753.
- [11] J. C. Quiroz, M. H. Yong, and E. Geangu. 2017. Emotion-recognition using smart watch accelerometer data: Preliminary findings. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 805–812.
- [12] J. Petersen, D. Austin, N. Mattek, and J. Kaye. 2015. Time out-of-home and cognitive, physical, and emotional well-being of older adults: A longitudinal mixed effects model. *PLoS One* 10, 10, Article e0139643 (2015).
- [13] D. J. Cook, M. Schmitter-Edgecombe, L. Jonsson, and A. V. Morant. 2018. Technology-enabled assessment of functional health. *IEEE Rev. Biomed. Eng.* 12 (2018), 319–332.
- [14] S. Majumder et al. 2017. Smart homes for elderly healthcare—Recent advances and research challenges. *Sensors*, 17, 11 (2017) 2496
- [15] A. Pratap, J. A. Anguera, and B. N. Renn. 2017. The feasibility of using smartphones to assess and remediate depression in Hispanic/Latino individuals nationally. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 854–860.
- [16] J. Austin, H. H. Dodge, T. Riley, P. G. Jacobs, S. Thielke, and J. Kaye. 2016. A smart-home system to unobtrusively and continuously assess loneliness in older adults. *IEEE J. Transl. Eng. Heal. Med.* 4, Article 7488979 (2016).
- [17] T. Grover and G. Mark. 2017. Digital footprints: Predicting personality from temporal patterns of technology use. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 41–44.
- [18] G. Sprint, D. Cook, D. Weeks, and V. Borisov. 2015. Predicting functional independence measure scores during rehabilitation with wearable inertial sensors. *IEEE Access* 3:1350–1366.
- [19] J. D. Howcroft, E. D. Lemaire, J. Kofman, and W. E. McIlroy. 2014. Analysis of dual-task elderly gait using wearable plantar-pressure insoles and accelerometer. In *Proceedings of the International Conference of the IEEE Engineering in Medicine and Biology Society*. 5003–5006.
- [20] M. Cella, L. Okruszek, M. Lawrence, V. Zarlenga, Z. He, and T. Wykes. 2017. Using wearable technology to detect the autonomic signature of illness severity in schizophrenia. *Schizophr. Res.* 195 (2017), 537–542.
- [21] D. J. Cook, G. Sprint, R. Fritz, and G. E. Duncan. 2018. Using smart city technology to make healthcare smarter. *Proc. IEEE* 106, 4 (2018), 708–722.
- [22] J. Wang, Z. Liu, and Y. Wu. 2013. Learning actionlet ensemble for 3D human action recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 36, 5 (2013), 914–927.
- [23] M. Mancini et al. 2016. Continuous monitoring of turning mobility and its association to falls and cognitive function: A pilot study. *J. Gerontol. Ser. A Biol. Sci. Med. Sci.* 71, 8 (2016), 1102–1108.
- [24] A. Ejupi, M. Brodie, S. R. Lord, J. Annegam, S. J. Redmond, and K. Delbaere. 2017. Wavelet-based sit-to-stand detection and assessment of fall risk in older people using a wearable pendant device. *IEEE Trans. Biomed. Eng.* 64, 7 (2017), 1602–1607.
- [25] D. J. Cook, P. Dawadi, and M. Schmitter-Edgecombe. 2015. Analyzing activity behavior and movement in a naturalistic environment using smart home techniques. *IEEE J. Biomed. Heal. Inform.* 19, 6 (2015), 1882–1892.
- [26] M. A. Ul Alam, N. Roy, S. Holmes, A. Gangopadhyay, and E. Galik. 2016. Automated functional and behavioral health assessment of older adults with dementia. In *Proceedings of the International Conference on Connected Health: Applications, Systems and Engineering Technologies*. 140–149.
- [27] S. Robben, G. Englebiene, and B. Krose. 2017. Delta features from ambient sensor data are good predictors of change in functional health. *IEEE J. Biomed. Heal. Inform.* 21, 4 (2017), 986–993.
- [28] A. Akl, B. Taati, and A. Mihailidis. 2015. Autonomous unobtrusive detection of mild cognitive impairment in older adults. *IEEE Trans. Biomed. Eng.* 62, 5 (2015), 1383–1394.
- [29] N. Dey, A. S. Ashour, F. Shi, S. J. Fong, and J. M. R. S. Tavares. 2018. Medical cyber-physical systems: A survey. *J. Med. Syst.* 42, 74 (2018).
- [30] Y. Zhang, M. Qiu, C.-W. Tsai, M. M. Hassan, and A. Alamri. 2015. Health-CPS: Healthcare cyber-physical system assisted by cloud and big data. *IEEE Syst. J.* 11, 1 (2015), 88–95.
- [31] H. M. Do, M. Pham, W. Sheng, D. Yang, and M. Liu. 2018. RiSH: A robot-integrated smart home for elderly care. *Rob. Auton. Syst* 101 (2018), 74–92.
- [32] R. H. Wang, A. Sudhama, M. Begum, and R. Huq. 2017. Robots to assist daily activities: Views of older adults with Alzheimer’s disease and their caregivers. *Int. Psychogeriatrics* 29, 1 (2017), 67–79.
- [33] M. Begum, R. Huq, R. H. Wang, and A. Mihailidis. 2015. Collaboration of an assistive robot and older adults with dementia. *Gerontechnology* 13, 4 (2015), 405–419.
- [34] T. L. Mitzner et al. 2010. Older adults talk technology: Technology usage and attitudes. *Comput. Human Behav.* 26, 6 (2010), 1710–1721.
- [35] G. Hoffman and W. Ju. 2014. Designing robots with movement in mind. *J. Human-Robot Interact.* 3, 1 (2014), 91–122.

- [36] M. Destephe, T. Maruyama, M. Zecca, K. Hashimoto, and A. Takanishi. 2013. Improving the human-robot interaction through emotive movements A special case: Walking. *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*.
- [37] A. Bisio et al. 2014. Motor contagion during human-human and human-robot interaction. *PLoS One* 9, 8 (2014), e106172.
- [38] A. Kupferberg, S. Glasauer, M. Huber, M. Rickert, A. Knoll, and T. Brandt. 2011. Biological movement increases acceptance of humanoid robots as human partners in motor interaction. *AI Soc.* 26, 4 (2011), 339–345.
- [39] A. H. Qureshi, Y. Nakamura, Y. Yoshikawa, and H. Ishiguro. 2016. Robot gains social intelligence through multi-modal deep reinforcement learning. In *Proceedings of the IEEE-RAS International Conference on Humanoid Robots*. 745–751.
- [40] Honda. 2007. Technical information. *Asimo*, 2007. Retrieved from <http://asimo.honda.com/downloads/pdf/asimo-technical-information.pdf>.
- [41] M. J. Mataric. 2017. Socially assistive robotics: Human augmentation versus automation. *Sci. Robot.* 2 (2017), 1–3.
- [42] K. M. Goher, N. Mansouri, and S. O. Fadlallah. 2017. Assessment of personal care and medical robots from older adults' perspective. *Robot. Biomimetics* 4, 1 (2017) 5.
- [43] Y.-T. Liao, C.-H. Zong, H.-H. Lee, and E. Tanaka. 2018. Development of kinect-based upper-limb assistance device for the motions of activities of daily living. In *Proceedings of the IEEE International Conference on Cyborg and Bionic Systems*.
- [44] Riken. 2015. The strong robot with the gentle touch. *News and Media*, 2015. Retrieved from http://www.riken.jp/en/pr/press/2015/20150223_2/.
- [45] J. Fasola and M. J. Mataric. 2013. A socially assistive robot exercise coach for the elderly. *J. Human-Robot Interact.* 2, 2 (2013), 3–32.
- [46] L. Pu, W. Moyle, C. Jones, and M. Todorovic. 2019. The effectiveness of social robots for older adults: A systematic review and meta-analysis of randomized controlled studies. *Gerontologist* 59, 1 (2019), e37–e51.
- [47] K. Wada, T. Shibata, T. Asada, and T. Musha. 2007. Robot therapy for prevention of dementia at home. *J. Robot. Mech.* 19, 6 (2007), 691.
- [48] P. Bovbel and G. Nejat. 2014. Casper: An assistive kitchen robot to promote aging in place. *J. Med. Device.* 8, 3 (2014).
- [49] J. Hoey, A. Von Bertoldi, T. Craig, P. Poupart, and A. Mihailidis. 2010. Automated handwashing assistance for persons with dementia using video and a partially observable Markov decision process. *Comput. Vis. Image Underst.* 114, 5 (2010), 503–519.
- [50] J. Hoey, T. Plotz, D. Jackson, A. Monk, C. Pham, and P. Olivier. 2011. Rapid specification and automated generation of prompting systems to assist people with dementia. *Pervasive Mob. Comput.* 7, 3 (2011), 299–318.
- [51] T. Kosch, P. W. Wozniak, E. Brady, and A. Schmidt. 2018. Smart kitchens for people with cognitive impairments: A qualitative study of design requirements. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 271.
- [52] S. Yan, K.-J. Lin, X. Zheng, and W. Zhang. 2019. Using latent knowledge to improve real-time activity recognition for smart IoT. *IEEE Trans. Knowl. Data Eng.* 2019.
- [53] J. Wang, Y. Chen, S. Hao, X. Peng, and L. Hu. 2019. Deep learning for sensor-based activity recognition: A survey. *Pattern Recog. Lett.* 119 (2019), 3–11.
- [54] B. Ghanem et al. 2018. ActivityNet large-scale activity recognition challenge 2018. Retrieved from <http://activity-net.org/challenges/2018/index.html>.
- [55] J. A. Stankovic. 2017. Research directions for cyber physical systems in wireless and mobile healthcare. *ACM Trans. Cyber Phys. Syst.* 1, 1 (2017), 1.
- [56] B. Das, D. J. Cook, N. Krishnan, and M. Schmitter-Edgecombe. 2016. One-class classification-based real-time activity error detection in smart homes. *IEEE J. Sel. Top. Signal Process.* 10, 5 (2016), 914–923.
- [57] M. F. Schwartz, L. J. Buxbaum, M. W. Montgomery, E. Fitzpatrick-DeSalme, T. Hart, and M. Ferraro. 1999. Naturalistic action production following right hemisphere stroke. *Neuropsychologia*, 37 (1999), 51–66.
- [58] M. M. N. Bienkiewicz, M.-L. Brandi, G. Goldenberg, C. M. L. Hughes, and J. Hermsdorfer. 2014. The tool in the brain: Apraxia in ADL. Behavioral and neurological correlates of apraxia in daily living. *Front. Psychol.* 5 (2014), 353.
- [59] S. Thrun, W. Burgard, and D. Fox. 2005. *Probabilistic Robotics*. The MIT Press.
- [60] G. Cheng, J. Han, P. Zhou, and D. Xu. 2019. Learning rotation-invariant and Fisher discriminative convolutional neural networks for object detection. *IEEE Trans. Image Proc.* 28, 1 (2019), 265–278.
- [61] J. Huang et al. 2016. Speed/accuracy trade-offs for modern convolutional object detectors. Retrieved from *CoRR*, abs/1611.1, 2016.
- [62] B. Reisburg et al. 2001. The Alzheimer's disease activities of daily living international scale (ADL-IS). *Int. Psychogeriatrics* 13, 2 (2001), 163–181.

- [63] M. Diel, M. Marsiske, A. Horgas, A. Rosenberg, J. Saczynski, and S. Willis. 2005. The revised observed tasks of daily living: A performance-based assessment of everyday problem solving in older adults. *J. Appl. Gerontechnology* 24 (2005), 211–230.
- [64] M. F. Folstein, S. E. Folstein, and P. R. McHugh. 1975. Mini-mental state. A practical method for grading the cognitive state of patients for the clinician. *J. Psych. Res.* 12, 4 (1975), 189–198.
- [65] Estimote. 2018. The physical world. *Software-defined*, 2018. Retrieved from <https://estimote.com/>.
- [66] G. Wilson et al. 2019. Robot-enabled support of daily activities in smart home environments. *Cogn. Syst. Res.* 54 (2019), 258–272.

Received May 2019; revised August 2019; accepted August 2019