

“Teach Me–Show Me”—End-User Personalization of a Smart Home and Companion Robot

Joe Saunders, Dag Sverre Syrdal, Kheng Lee Koay, Nathan Burke, and Kerstin Dautenhahn, *Senior Member, IEEE*

Abstract—Care issues and costs associated with an increasing elderly population are becoming a major concern for many countries. The use of assistive robots in “smart-home” environments has been suggested as a possible partial solution to these concerns. A challenge is the personalization of the robot to meet the changing needs of the elderly person over time. One approach is to allow the elderly person, or their carers or relatives, to make the robot learn activities in the smart home and teach it to carry out behaviors in response to these activities. The overriding premise being that such teaching is both intuitive and “nontechnical.” To evaluate these issues, a commercially available autonomous robot has been deployed in a fully sensorized but otherwise ordinary suburban house. We describe the design approach to the teaching, learning, robot, and smart home systems as an integrated unit and present results from an evaluation of the teaching component with 20 participants and a preliminary evaluation of the learning component with three participants in a human–robot interaction experiment. Participants reported findings using a system usability scale and ad-hoc Likert questionnaires. Results indicated that participants thought that this approach to robot personalization was easy to use, useful, and that they would be capable of using it in real-life situations both for themselves and for others.

Index Terms—Activity recognition, robot companion, robot learning, robot personalization, robot teaching.

I. INTRODUCTION

ASSISTIVE robots in “smart-home” environments have been suggested as a possible cost and care solution to demographics changes characterized by an increasing elderly population [1], [2]. The vision is that service robots are available in the home to help and assist elderly residents. Furthermore, the robot might also motivate and provide active support in terms of *reablement*—defined as “Support people ‘to do’ rather than ‘doing to / for people’” [3]—and *co-learning*—working together to achieve a particular goal. Thus, the assistive robot and the person form a partnership which is ever changing and evolving to meet the changing needs of the elderly person as they age, the robot effectively becoming a trusted companion to the person. We define this mechanism of providing support, assistance, and active engagement over time as *personalization*. This paper describes an approach to service robot personalization based on

Manuscript received February 18, 2015; revised May 19, 2015; accepted May 25, 2015. Date of publication November 9, 2015; date of current version January 20, 2016. This work was supported in part by the European project ACCOMPANY (Acceptable robotics COMPanions for AgeiNg Years) under Grant 287624. This paper was recommended by Associate Editor D. B. Kaber.

The authors are with the Adaptive Systems Research Group, University of Hertfordshire, Hertfordshire AL10 9AB, U.K. (e-mail: j.l.saunders@herts.ac.uk; d.s.syrdal@herts.ac.uk; k.l.koay@herts.ac.uk; n.burke@herts.ac.uk; k.dautenhahn@herts.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/THMS.2015.2445105

end-user *robot teaching and learning* designed to be used by carers, relatives, and elderly persons themselves. Personalization has been shown in longitudinal studies to reinforce rapport, cooperation, and engagement with a robot [4].

The work described in this paper uses a commercially available robot, the Care-O-bot3 [5]. The robot resides in a fully sensorized but otherwise completely standard British three bedroom semi-detached house (we call this the *robot house*). This environment is more ecologically valid than a scientific laboratory for evaluating human–robot interaction (HRI) studies.

II. PROBLEM DEFINITION

A. Co-learning and Reablement

The idea of co-learning in this context refers to the situation whereby a human user and a robot work together to achieve a particular goal. Typically, the robot can provide help and assistance, but in return also requires help and assistance. Usually, the human teaches the robot how to solve a problem; however, the robot can also assist by suggesting to the human that it has particular capabilities and techniques which may prove fruitful. This concept is extended by considering that the robot will need to learn from the user about the user’s activities and subsequently be able to exploit this information in future teaching episodes. This means that cooperation will typify the user’s interaction with the robot. The concept of reablement [6] exploits the co-learning capability in order not to disenfranchise the human partner. Thus, rather than passively accepting imposed solutions to a particular need, the user actively participates in formulating with the robot their own solutions and thus remains dominant of the technology and is empowered, physically, cognitively, and socially. This idea is extended by ensuring that the robot engages in empathic and socially interactive behavior. For example, the robot should not attempt to encourage immobility or passivity in the user, but to reable the user by making motivating suggestions to persuade the user to be active or engage in an activity in the home. For example, it could prompt the user to carry out tasks, for example, writing a greeting card after reminding the user of a relative’s birthday, or bring relevant events to the user’s attention and suggest to the user an activity in order to avoid social isolation. Thus, the user–robot relationship is one of mutually beneficial support, assistance, and companionship.

B. Background

Achieving personalization presents many challenges for a companion robot. Simple scripting of interactions will not achieve the above aims due to the dynamics of the interaction and the key requirement of the robot to develop and learn.

1) *Robot Teaching and Learning*: Approaches in robot learning attempt to derive a policy based on one or more demonstrations from the user and subsequently execute that policy at the appropriate time. Other key challenges are how to refine the policy from further demonstrations, how to scaffold different policies together to form more complex robot behaviors, and how to allow the robot to inform the user of its existing repertoire of policies. For a more detailed surveys, see [7] and [8].

2) *Learning by “Following”*: This approach typically involves both robot and user sharing a close context. The robot often uses a vision system or some other sensory modality (e.g., infrared sensors, electronic markers) to detect the presence of the user and then follow him/her. By closely following the user, the robot is able to approximately experience the same perceptions as those experienced by the user. Thus, the “state” of the user, which is not normally perceivable by the robot, can be perceived indirectly. The “following” approach has been used in work by Nicolescu and Mataric [9], where a mobile robot tracks a human teacher’s movements by following the teacher and matching predicted postconditions against the robot’s current proprioceptive state. It then builds a hierarchical behavior-based network based on “Strips” style production rules [10]. This work attempts to provide a natural interface between robot and a teacher (who provides feedback cues) whilst automatically constructing an appropriate action-selection framework for the robot. During the learning process the robot can use external environmental perceptions and any available internal proprioceptive feedback in order to replicate the user’s behavior (for a more detailed discussion, see [11]). In the personalization research reported in the current paper, we exploit many of these techniques; however, external sensory cues are provided to the robot exclusively via the smart home sensors.

3) *Behavioral Cloning*: Behavioral cloning is used primarily as a way of encoding human knowledge in a form that can be used by a computational system. The actions of a human subject, who will be typically operating a complex control system, are recorded and analyzed. The actions and decisions are extracted and used to control the system without human presence. An example of behavioral cloning is Claude Sammut’s “learning-to-fly” application [12], [13], where recordings of control parameters in a flight simulator flown by human subjects were analyzed using Quinlan’s C4.5 induction algorithm [14]. The algorithm extracts a set of “IF–THEN” control rules. Van Lent and Laird extended this work by providing a user interface which could be marked with goal transition information [15]. This allowed an action selection architecture to be constructed using “Strips” style production rules [10]. In the current paper, we also provide the house resident with an interface for teaching robot behaviors based on previously learnt activities using Quinlan’s C4.5 rule induction system. The resulting robot behavioral rules are also based on a production rule approach.

4) *Learning by Demonstration*: Learning by Demonstration normally refers to the direct interaction between a human teacher and a robot.¹ The interaction is direct because the teacher sends

instructions to the robot directly through some external control mechanism (e.g., a joystick or screen based GUI). This direct approach avoids many of the complexities of the *Correspondence Problem* [16]. Early work by Levas and Selfridge [17] controlled a robot via teleoperation and then used the robot’s proprioceptive feedback to construct a set of production rules. Teaching service robots by observing humans in this manner have also been carried out by Dillman *et al.* [18]–[21]. Kaiser trained various robotic platforms in order to compute a control policy using function approximation techniques and recognized the important role of the human teacher in providing feedback. Similar observations have also been made by Thomaz and Breazeal [22].

5) *Learning From Observation*: Learning from Observation normally decreases the closeness of shared context between learner and user. Thus, the robot operates by sharing context with the user but at a distance. This research relies on recognizing human motions and thus faces a difficult vision problem. In order to obviate this problem, complex vision techniques are sometimes employed. Often, however, the problem is simplified using colored markers or some other tagging technique. Examples of learning using an observational approach include [23] and [24] where hierarchical and symbolic representations of assembly tasks are learned from human demonstration. Johnson and Demiris [25] use learning from observation in their work where coupled inverse and forward models [26], [27] are used to allow a robot to imitate observed human actions and recognize new actions. In the smart home context described in this paper, we do not directly use observational approaches but use the human feedback derived from the house sensor activations (including human location tracking).

C. Teaching and Learning in Smart Home Environments

Teaching, learning, and adaptation in smart home environments tend to be based more on automatic service discovery where the home automatically learns the daily activities of the resident. Often called “Cognitive Robotic Ecologies” they attempt to understand the requirements of the house residents based on perception, planning, and learning from the house “ecology” and derive robotic actions to service these requirements. These methods face problems in identifying the information needed to make these judgments and to identify the appropriate teaching information to adapt such services.

Typical approaches include capturing and merging sensor information via machine learning techniques and then predicting resident behavior [28]–[30], the majority of which use labeled training examples built from annotation of resident activities. However, labeling can be costly and time consuming.

III. METHODOLOGICAL FORMULATION

Our work allows the house resident to personalize the robot to meet their changing needs and to exploit the robot’s existing competencies to achieve this where necessary. All basic activities, be they robot behaviors or house sensory states, can be easily interpreted by the house resident. Furthermore, the underlying design ensures that any new behaviors or activities can be interpreted as basic activities and exploit any services that apply to these activities.

¹“Learning by Demonstration” is often also used in a wider sense to denote all of the research areas that study robot teaching.

A. Extending the Idea of a Sensor

Consider a situation where the elderly resident has a robot that is capable of navigating autonomously around the house, can move to the user’s location, is equipped with a raisable tray, and has the ability to “speak” text strings. She would like the robot to always be present in the kitchen when she is using the microwave in order to carry items back to the dining table. She might teach the robot to do this by providing simple directives such as

```
If the microwave is on
then go to the kitchen and raise
    your tray
```

In this example, the microwave sensor is a basic physical sensor, and the robot actions are navigation and tray actuation. Simple sensory information could also be enhanced with temporal constraints. For example:

```
If the microwave has been on
    for more than 5 minutes
Then go to the user location
    and say 'the microwave is
    still on'
```

Furthermore, the simple sensory states could be replaced by states with higher levels of meaning. For example:

```
If 'food is being prepared'
then come to the kitchen
```

where the sensory state “food is being prepared” is derived from activity recognition (for example recognizing that the microwave or main oven or fridge were being used). Additionally, these higher states could be temporally extended:

```
If 'food is being prepared'
    and this has been happening
        for more than
        30 minutes
then go to the user location
    and say 'I think you are making
    a meal, do you need help?'
```

Similar grouping of basic robot actions should also be possible. Thus, simple sets of robot actions such as

```
go to user location, lowering tray
if tray empty
```

could simply be labeled:

```
come to me
```

By enabling constructs of this kind, the robot behavior personalization is enhanced. Consider a carer setting up a robot behavior to remind the elderly resident that her daughter visits her in the afternoon on Tuesdays and Thursdays:

```
If it is Tuesday or Thursday and 1pm
then 'come to me'
    and say 'Irene is coming to
    visit you today'
```

In this paper, the definition of “sensory state” is expanded to include both physical and semantic states and the definition of robot action expanded via the ability to “scaffold” robot behaviors [31] to create more complex but semantically simpler behaviors. This makes the overall system easier to understand and hides the technical complexities of robotics and smart home systems from the user.

Two complimentary approaches to achieving this level of personalization were designed and called “Teach Me/Show Me.” These were implemented as a program running on a laptop computer. The “Teach Me” system allows residents’ to define and test robot behaviors based on both house sensory activities and basic robot actions. These new behaviors, once defined, can be used to create more complex behaviors. The “Show Me” system allows the resident to “show” the robot new activities (such as “preparing a meal”) by simply carrying out that task. Once learned that activity becomes part of the available sensory activities exploitable by the Teach Me system.

An advantage of this approach is that there is no pre-labeling of activities. Labeling of sensory combinations of all types is effectively carried out by the resident. The resident thus personalizes requirements and is thus enabled and enfranchised by being at the center of the personalization process.

B. System Architecture

We regard the house as one entity rather than as a collection of individual parts. In practice, this means that the house sensor information is considered to be no different from robot sensor information, the sensory information derived from the occupants activities or from semantic sensors. This provides the bedrock for the main focus of our work enabling co-learning and reablement by not artificially treating the robot, user or the house as separate entities but rather focus the generation of behavioral activity on the complete system.

1) *Robot House Ontology:* The robot house consists of sensors, locations, objects, people, the robot and (robot) behaviors. These were analyzed to yield a house ontology instantiated in a “mysql” database. Procedural memory, defined as the robot actions together with pre- and post-behavioral conditions, is also held as tables in the database. However, the rules are encoded as SQL statements, which refer back to the semantic information created by the sensor system.

2) *Robot Capabilities:* We use the Care-O-bot3 (see Fig. 2) designed for research in assistive environments. It uses ROS navigation [32] using its laser range-finders to update a map of the house in real time and can thus navigate to any given location whilst avoiding obstacles and replanning routes. The robot is equipped with facilities for manipulating the arm, torso, “eyes,” robot LED’s, tray and has a voice synthesizer to express given text. High-level commands are sent via the ROS “script server” mechanism and interpreted into low-level commands by the robot software. Typical commands would be “raise tray,” “nod,” “look forward,” “move to location x,” “grab object on tray,” “put object x at location y,” and “say hello.”

3) *Robot House Sensors:* All sensory information (both from physical and from semantic sensors) is held in a “sensors” table and a “sensor logging” table in the database. Each



Fig. 1. Interior of the robot house “living room” with the Care-O-Bot3 robot. The images on the left side of the picture show the view from the robot camera, the robot itself, and the real-time mapping visualization showing person and robot location.

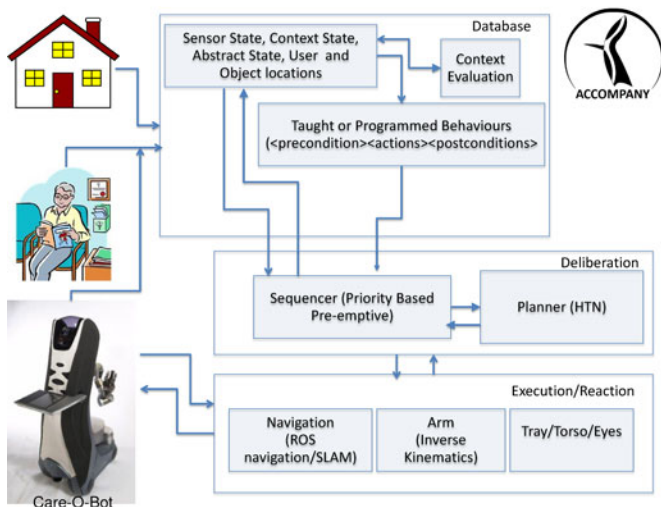


Fig. 2. Layers in operation in the robot house. Sensory information from the robot, house and people together with semantic sensors update the database in real time. Taught behaviors use these sensors to access behavioral preconditions and may set semantic sensors during execution. All behavioral preconditions are continually evaluated by the scheduling system and become available for execution if all preconditions are met. Actions that require planning call an HTN planner. Lower level functions such as navigation and arm manipulation work at a reactive level.

individual sensor is held as a row in the sensors table and each row provides the instantaneous value of the sensor as well as the time it changed and its previous value. Each row in the logging table contains the historical sensor value over time. The “TeachMe” system uses only the current and previous sensor values, whereas the “ShowMe” system exploits the historical sensor log. The robot house (see Fig. 1) contains around 50 “low level” sensors. These range from electrical (e.g., fridge door open), to furniture (e.g., cupboard door closed), to services (e.g., toilet flushing, taps running) and pressure devices (e.g., bed occupied). Sensory information from the robot is also sent to the database or, for high throughput, is acquired via ROS messaging [32]. In addition user locations are known to the robot via ceiling mounted cameras [33] and robot locations are available via ROS navigation [32] in a common framework. There are an

unlimited number of semantic sensors dependent on what the resident teaches the robot.

4) *Behavior Encoding*: Behaviors are automatically generated from the teaching facilities in Section III-C. However, each behavior generated follows a template similar to Nilsson’s T-R formalism [34] of evaluating preconditions, followed by execution of robot actions and updating of postconditions. Preconditions can be applied to any form of sensory information, both set by the environment or set at a “semantic” level. An example of such a behavior:

```

IF    the oven has been on for 90 minutes
      // house sensor precondition
AND  the user has not already been
      reminded
      // semantic sensor precondition
THEN 'come to me'
      // scaffolded robot action
      say 'The oven has been on for
          a long time'
      // basic robot action
      update the database to signal
      that the user has been reminded
      // set semantic sensor
      // post-condition

```

The preconditions would be automatically encoded by the teaching system as SQL statements (two SQL statements representing preconditions would be generated for the example given above):

```

SELECT * FROM Sensors
WHERE  sensorId = 50
      AND value = 'On'
      AND lastUpdate+INTERVAL 5400
      SECOND <= NOW()
SELECT * FROM Sensors
WHERE  sensorId = 701
      AND value = 'notReminded'

```

If a row is returned from the execution of the SQL statement, then that precondition is deemed to be true, otherwise false. Typical robot actions, e.g., calling the navigation system to move the base, making the robot say something, and updating a semantic sensor are shown below:

```

base,0,[4.329:0.837:51],999,wait
speak,0,The oven has been on for a
long time
cond,701,reminded

```

These commands, depending on the command type (e.g., for the example above, “base” moves the robot, “speak” invokes the voice synthesizer, and “cond” sets the value of a semantic sensor), would then either be sent to the planner (see Section III-B9), or sent directly to a lower level control module if planning was not required.

5) *Sensors and Sensor Abstraction*: All sensory information updates the database in real time and all robot behaviors

continually retrieve information from these sensors to assess whether their behavioral preconditions are met allowing behavioral scheduling and execution (explained in Section III-B11). Behaviors will continue to execute if their preconditions remain true or unless they are pre-empted by a higher priority behavior.

6) *Semantic Sensors*: In order to cope with ongoing events in the house which are not reflected by the physical house sensors a set of “semantic sensors” can be created by the teaching system, e.g., a sensor with the label “User has been reminded that the oven is on.” This latter sensor would be set to ‘reminded’ following the spoken oven reminder in the example in Section III-B4 above. Similarly, an activity context recognition system can update semantic sensors in real time based on the “Show Me” system described in Section III-C2 below. Thus, if the user has shown the system what activities constitute “preparing a meal,” then the “preparing a meal” semantic sensor would be set to “true” when these events occur.

7) *Temporal Aspects of Sensors*: Using sensors at a physical and semantic level provides the opportunity to apply temporal constraints. Consider a doorbell; this type of sensor is “on” only for a short period of time, and thus, rather than ask “Is the doorbell ringing?” we would ask “has the doorbell rung within the last 30 s?” This is checked by holding episodic values and we can query previous values at a previous point in time:

```
SELECT * FROM Sensors
WHERE     sensorId = 59
        AND lastActiveValue > 0
        AND     lastUpdate+INTERVAL 30
              SECOND >= NOW()
```

The further capability of assessing how long a sensor has been active (or inactive) allows for greater behavioral expressivity [35]. For example, “Has the user been sitting on the sofa for longer than 2 h?” “has the user been reminded to call his friend Albert this week?” These encoding facilities can, therefore, cope with a very wide range of situations and capture information related to current activity, past activity, and socially desirable activity, the latter being primarily set through the creation of semantic sensors.

8) *External Sensors and External Actions*: The sensor system provides a standardized way of encoding information and provides possibilities for associating semantic sensors with other, typically external, events. For example, by polling an external weather service it would be possible to set a “weather” sensor. This could be checked by a behavior which might suggest that this was a good day for a walk, or to do some gardening. This way, the idea of *reablement* can be operationalized. External actions could also be run, for example calling a text messaging (SMS) service. For example, a behavior that checks whether the bed pressure sensor had been active for more than 12 h and that there had been no activity in the kitchen might then send a text message to the user’s caregivers suggesting that the person might need assistance to get out of bed.

9) *Planning*: Our general approach is to plan only when needed and when necessary. The overall behavior of the system is driven primarily by the environmental conditions via house or semantic sensors values queried via behavioral preconditions.

Behaviors are explicitly scheduled. However, there are instances where, due to multiple choices being available for robot action (e.g., in a multiroom environment, navigation may take multiple paths), or when there is conflict between available resources when planning is necessary. We consider that creating planning domains to be too complex for end-user involvement and we precode these where necessary.

10) *Planning Domain*: We use an open-source state-of-the-art hierarchical task network (HTN) planner (SHOP2 [36]) to cope with these situations. We follow the approach in [37] and [38], in that each planning domain is individually coded in the lisp-like syntax of SHOP2 and called when the high level action is required. SHOP2 returns the planning actions as robot actions. After each action execution, we recall the planning component just in case the environment has changed between actions.

11) *Preemptive Scheduling*: Behaviors can be created via a technical interface [39], used when the system is first installed by technical personnel or by the end user using the “TeachMe” facility described in this paper. The “technical” interface allows a priority to be given to each behavior whereas the “TeachMe” system sets all created behaviors to have the same priority. On execution, the scheduling system continually checks all of the preconditions of all of the behaviors (in a manner similar to [40]). Should all of the preconditions of a behavior be satisfied the behavior becomes available for execution, with the highest priority behavior being executed first. Priority ties result in a random choice of behavior for execution. Due to continual checking of all behavioral preconditions, behaviors may become valid or invalid for execution as the currently executing behavior operates. In this manner, the set of environment and semantic sensors drive behavior execution. Some behaviors can also be set as non-interruptible, for example if a behavior was reporting on a critical house event—such as the bathroom taps running for a long time.

C. Teaching and Learning Interfaces

The teaching interface allows users to create robot behaviors, the learning interface allows users to create higher level semantic sensors for use by the teaching system. For example, the user might create a sensor called “relaxing in the afternoon” using the learning system and exploit it in a robot behavior such as “If I am ‘relaxing in the afternoon’ for longer than 3 h remind me to take some exercise.”

1) *Teaching Interface—“Teach Me”*: In order to create behaviors, the user as a minimum would need to specify *what needs to happen* (the actions of the robot) and *when those actions should take place* (setting preconditions based on the values of physical or semantic sensors). Having specified “what” and “when” the system automatically generates many of the sub-behaviors required to operationalize the system. It does this by using templates. This simplification trades generality for ease of use so that the system can be used by non-experts in real-life scenarios.

Consider a user who wants to be reminded to take medicine at 5 P.M. If we were to create this task, individual behaviors would need to be created to associate each precondition with the

TABLE I
TAUGHT ROBOT BEHAVIORS INCREASING IN BEHAVIORAL COMPLEXITY

<i>Taught Behaviors - Set 1</i>
Whenever you open the microwave oven, make the robot come to the kitchen and wait outside. If the TV is ON, and you are sitting on the sofa, make the robot join you at the sofa. If the doorbell rings and the TV is ON, make the robot say "There's someone at the door" and then go to the Hallway.
<i>Taught Behaviors - Set 2</i>
Make the robot come to the table and remind you to always call your friend "Mary" on Thursday at 2 P.M. On Mondays, Wednesdays, and Fridays make the robot stand in the hall and remind you that your lunch is being delivered at 12:30 P.M. If you are sitting on the sofa for more than 30 min, make the robot come to the sofa and tell you to take some exercise. Make the robot do this again after another 30 min if you are still sitting there. Make the robot come to the table and remind you to take your medicine at 12 P.M. and 4 P.M. every day, yellow pills at 12, pink at 4 P.M.
<i>Taught Behaviors - Set 3</i>
Make the robot come to the sofa and tell you to "move about a bit," if, in the afternoon, you have sat on the sofa for longer than 1 h continuously. If it is after 9 P.M., and you have left the sofa for over 5 min and the TV is ON, make the robot go to the hall entrance and say "turn off the TV." If the microwave has been open or on for more than 1 min, make the robot come to the table and tell you that the microwave is open or ON. Make the robot remind you every minute until the microwave is turned OFF and door is closed.

appropriate sensor, including the semantic sensors, effectively creating two behaviors, one to carry out the task and one to reset conditions, as follows:

- 1) The first behavior would need to check that the time was after 5 P.M. and that the user had not been already reminded, i.e., a "user not yet reminded" semantic sensor would be true. If both of these conditions are true, then the robot carries out a procedure of moving to the user and saying "It's time for your medicine" then resetting the semantic sensor to false to indicate that the user has been reminded.
- 2) At some point later, a second behavior would need to run, which in this example would be: if after midnight, reset the "user not yet reminded" sensor to true so that it can fire the next day.

Thus two behaviors need to be created, and careful alignment of reminder rules need to be inserted.

However, the sort of behaviors (see Table I) that we envisage users setting up themselves tend to follow a set of common templates, e.g., diary like functions, or direct actions based on sensory conditions in the house. We can, therefore, exploit these templates to generate the appropriate conditional logic. The template in the example above is based on "diary" like conditions and the automatic setting and creation of support behaviors (such as the resetting behavior above). In this manner, much of the cognitive load is removed from the user and left to the behavior generation system. Co-learning is operationalized by allowing the robot to provide details of its existing sets of skills that can then be exploited by the user. Reablement is supported simply in the act of teaching the robot.

The standard template for "diary like" robot actions is as follows:

Entered by user via GUI:

```
reminderTime = t (e.g. 5pm)
textItem e.g. 'Have you taken
your medicine?'
repeatAfter = n (e.g. 60 seconds)
<other robot actions> e.g. ''Move
to user''
```

Created automatically:

```
Cond-Reminder = TRUE
Cond-Remind-again = FALSE
```

Then create the following robot behaviors automatically:

- 1) ReminderX-reset: % resets conditions

```
IF NOW between midnight and t
AND
Cond-Reminder = FALSE
SET Cond-Reminder = TRUE
SET Cond-Remind-again = FALSE
```

- 2) ReminderX: % the actual diary reminded

```
IF NOW >= t
AND
Cond-Reminder = TRUE
EXECUTE <other robot actions>
SAY <text item>
SET Cond-Reminder = FALSE
SET Cond-Remind-again = TRUE
```

An example of the user teaching GUI is in Figs. 3 and 4 and displays the actions a person would use to create the example behavior above. The steps consists of "what" the robot should do followed by "when" the robot should do it.

- 1) The user chooses to send the robot to the current user location and then presses a "learn it" button. This puts the command into the robot memory.
- 2) Then the user makes the robot say "It's time for your medicine." This is not in the robot's current set of skills and is entered as a text input. This is followed by a press of the "learn it" button.
- 3) Now the two actions are in the robot's memory and the user completes the "what" phase and starts on the "when" phase.
- 4) The user is offered a number of choices including reacting to events in the house, or user or robot locations or a diary function (second screen in Fig. 3).
- 5) The user chooses a diary function and enters 17:00 in the "at this time" box (first screenshot shown in Fig. 4).
- 6) Again this is followed by pressing the "learn it" button.

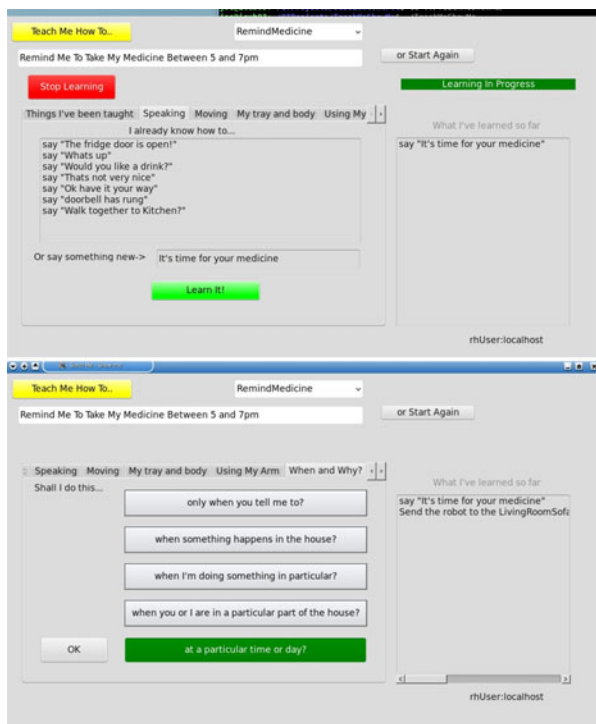


Fig. 3. Screenshots of the teaching interface (note that not all screens are shown—see main text and Fig. 4). In the top figure, the user has entered the words that the robot is meant to say. The second screen allows choice of differing activities, such as polling sensors or setting diary events.

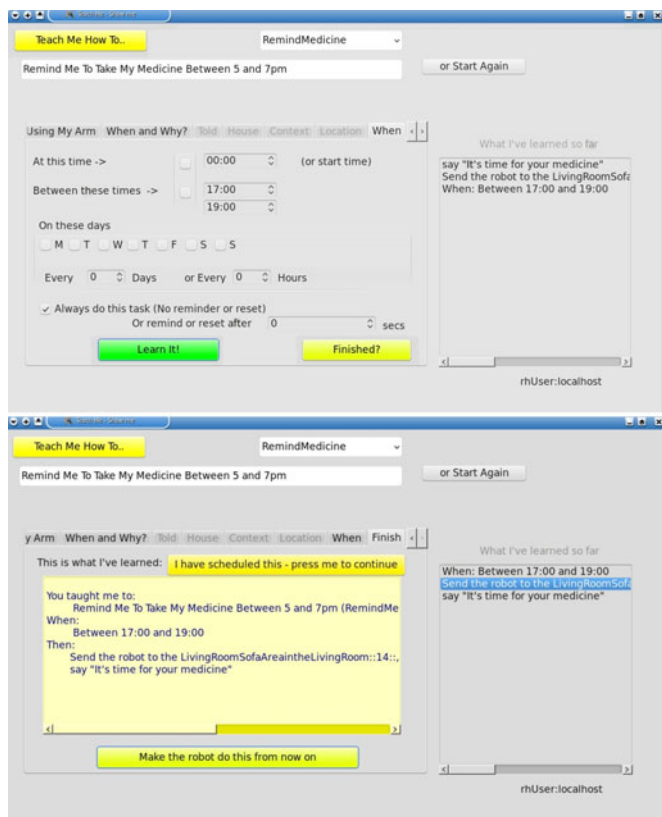


Fig. 4. Final 2 screenshots of the teaching interface. The top image shows the diary option selected in this case and the condition “after 5 P.M.” is entered. The bottom screen shows the final behavior created.

7) Having completed both “what” and “when” phases the user is shown the complete behavior for review and can modify it if necessary (bottom of Fig. 4).

8) Once happy the user presses a “make me do this from now on” button and the complete behavior becomes part of the robot behavioral repertoire.

2) *Learning Interface—“Show Me”*: The “Show Me” approach is contingent on the house occupant indicating to the robot that activities are underway. For example, the person might indicate that he or she are “preparing food.” Activities typically have a nested and sometimes hierarchical nature. For example, “preparing food” might include “making a hot or cold drink” or “using the toaster.” The start and end times and durations of the main task and the sub-tasks are variable. However, when any of the subtasks are active (e.g., using toaster) the main task must also be active (i.e., preparing food).

Consider that the person has indicated to the robot that he or she are “preparing food” and at some point also indicated that he or she are now “using the toaster.” If the robot learns the set of sensory activities associated with these tasks it should be able to recognize them when they occur in the future. Thus, the robot would recognize when the toaster is active and infer not only that “using the toaster” is true but also that “preparing food” is true.

Given that these activities can be recognized by the robot (via the house sensory system), it would then be possible to exploit these in the teaching system and the person would be able to teach the robot based on the higher level semantics associated with the task. For example, the user might teach “When I am ‘Preparing food,’ the robot should come to the kitchen and raise its tray.”

The learning system provides symbolic entries by automatically creating semantic sensors labeled with the descriptive term (e.g., “preparing food”) provided by the user. These can then be exploited to create new behaviors on the robot.

The challenges for a learning system are to recognize that learnt situations can be active in parallel, have an implicit nested hierarchy, and that higher levels in the hierarchy (typically) represent higher level of semantic knowledge. These need to be represented as lexical symbols in the memory architecture which the teacher can then exploit.

3) *Approach to Learning*: To learn typical activities in the house the robot needs to recognize when these situations re-occur. This recognition would be primarily based on the current sensory state of the house; however, in more complex circumstances, both the historical sensory state and a predicted future sensory state may also be necessary (for example, in historical terms, to recognize that the postman called this afternoon, or in the predicted sense, that the house occupant is likely soon to go to bed). In the work presented in this paper, we only consider the current sensory state.

We also have to consider that the certainty of situations cannot always be represented by a simple true/false dichotomy. For example, if I am in the kitchen it is likely I am preparing food, but it is not a certainty. The confidence of the task assessment by the robot has to be considered.

Our approach falls under the banner of ambient activity recognition in that house resident activities are modeled by analyzing a sequence of ambient sensor events. The various approaches to this research area typically apply supervised machine learning techniques such as decision trees/rule induction [41]) (as is used in the studies presented in this document), HMM's and dynamic Bayesian networks [42], template matching techniques such as k-NN [43] or dynamic windowing techniques [29]. Sensor data are typically pre-labeled by an external observer (although some techniques also search for common patterns in daily activities [28]). Our approach differs from a strict supervised learning approach in that the house resident is responsible for "labeling" the data and does this by providing the label and then carrying out the activity, while the system records and automatically assigns the label to the sensory data. The newly acquired activity can be subsequently used for direct robot teaching. Activity recognition is based on streaming vectorized sensor data—an approach which allows multiple activity patterns to be recognized in parallel.

The memory system is based on rule sets held as behaviors; these are human readable and taught by the human using the teaching system. Ideally, a learning system should be human readable. We employ a rule induction approach to learning based on Quinlan's C4.5 Rule induction algorithm (C5.0) [14] which allows generation of rule sets in human readable form.

4) *Verification of Approach:* In order to verify the plausibility of our approach, we exploited some existing end user behavior data [44].

In these previous studies, 14 participants were asked to carry out a series of typical daily activities within the house. Each participant took part in two sessions of approximately 45-min duration each. In the first (training) session, the experimenter suggested to the participant particular tasks that should be carried out. In the second (test) session, the experimenter asked the participant to carry out any of the tasks (or none) at their own discretion. All house sensory data were recorded and all sessions videotaped. The video tapes were annotated by Duque *et al.* [44] and an external observer and marked with the task and subtask start and end times. These were then matched against the recorded sensory data. Duque *et al.*'s. [44] aim was to *manually* create a rule-set, derived from the training sessions, which could be applied to the test data and accurately predict the activity that was being carried out by the participant. This rule set was constructed and applied to the test data resulting in recognition accuracy (based on precision, recall and accuracy) of over 80%. In the work presented in this paper, we tested the plausibility of our approach by replacing the designer generated rules with rules *automatically derived* using the C5.0 algorithm. We then assessed the performance of this approach.

The training data for the 14 participants was used to train a learner using C5.0 with boosting over 10 trials. The learner was then applied to the test data and the resulting performance analyzed for four activity states displayed on ROC curves (see Fig. 5).

Each data point in the ROC curves indicate a participant's training or testing session. Also shown are the combined results after aggregation of data of all of the 14 participants into




Fig. 5. ROC curves showing the results of the applying the induction system on both the training and testing data for four categories of classification.

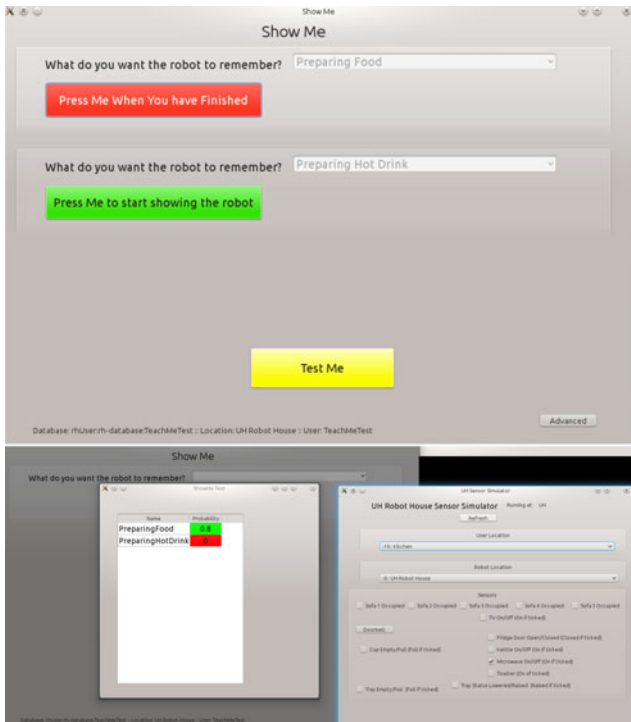


Fig. 6. “Show Me” learning GUI. Here, the user has entered “Preparing Food” and when ready presses the “press me to start showing the robot” button. They then carry out actions associated with preparing food (e.g., starting the microwave oven). If a subtask is required (in this case “Preparing a Hot Drink”), the user can continue to enter new tasks up to a maximum of three levels deep. Once each task completes, the user presses the red “Press me when you have finished” button. Testing can be carried out by pressing the “Test Me” button. This operates in real time and allows the user, while repeating the task, to check if the system correctly identifies it. A probability % is also given based on the predicted accuracy of the real-time classifier using the learned rules. The color of the classifier symbol turns green if the probability exceeds 50%. Note that the system automatically creates lexical symbols which are then available within the robot teaching interface. In the testing example (shown being tested with the house simulator), the microwave is ON; therefore, the system infers that “Preparing Food” is 80% certain. However, as the kettle is off, “Preparing Hot Drink” is very unlikely (0%).

one dataset. Clusters that occur in the top left quadrant of the ROC curves indicate a strong level of learning and recognition performance.

The ROC analysis indicated that such a learning approach can allow the robot to recognize human activities in the robot house. However, in a “real” situation we are faced with having no observer of human actions and no annotator of those actions to derive a classification set. To address this issue, we allowed the *house occupant to become the observer/annotator* by informing the robot when tasks are starting and finishing. To carry this out, an end-user training GUI was developed which we called “Show Me” (see Fig. 6). The GUI allowed users to state what they are currently doing (up to three hierarchical levels) and subsequently test whether the system correctly recognizes these actions.

5) *Learning and Execution Mechanism*: Data for the induction algorithm are held as a table of single row sensor vectors each labeled with the user-defined text provided by the GUI. The sensor vectors are used by C5.0 to produce its rule sets. These rule sets are then applied in real time to incoming sensory data from the house. The effectiveness of the rule set is

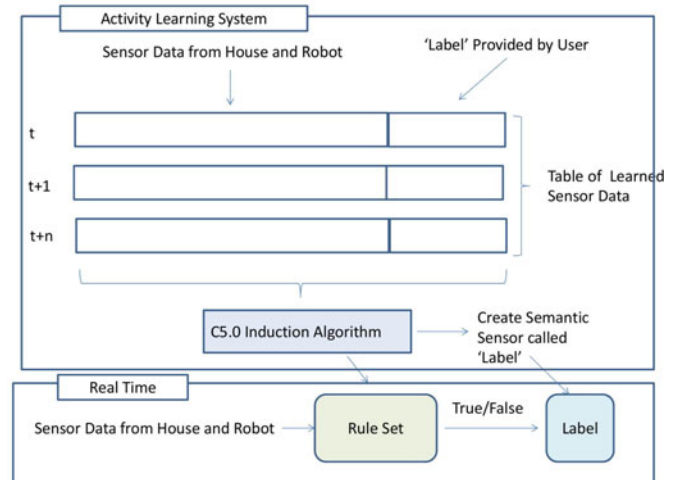


Fig. 7. “Show Me” system first asks the user to provide a label for the activity and captures vectorized sensor data to a file in real time. The recorded file is subsequently processed by the C5.0 algorithm resulting a rule set. The system also creates a semantic sensor labeled with the name given by the user. Real-time sensory data from the house and robot is queried by the rule set generated by C5.0, which results in the labeled semantic sensor being set either true or false.

expressed by C5.0 as a percentage. If this percentage exceeds 50%, the labeled semantic sensor is set to true, otherwise false. A pictorial representation of the process is shown in Fig. 7.

IV. EVALUATION OF THE TEACHING AND LEARNING SYSTEMS

A. Procedure for the Teaching System—“Teach Me”

The evaluation of the template-based teaching system involved 20 participants. The experimental procedure is outlined in Table II.

Each participant was introduced to the experimenter, a technician, and the experiment psychologist. The technician was present to ensure the safety of the participant (a requirement of the ethics agreement) and stationed in a part of the room outside the main interaction area.

The psychologist asked the participant to complete: a consent form, a demographics form, a questionnaire assessing computer and robot experience, and the Ten Item Personality Inventory (TIPI) [45].

The psychologist retired to a different room. The experimenter then explained the purpose of the experiment, the nature of the sensorized house and the capabilities of the robot (in this experiment, the robot capabilities were restricted to moving to differing locations and speaking, although the tray and arm/gripper were visible).

The robot had previously been taught to approach the experimenter and participant and to introduce itself by saying “welcome to the robot house.” This gave the experimenter a chance to explain the robot capabilities and for the participant to see the robot in action.

Examples of three sets of behaviors, each with increasing complexity, were shown to participants (see Table I). The behavior relating to “answering the doorbell” in set 1 was used by the experimenter to show the participant how to use the teaching GUI.

TABLE II
EXPERIMENTAL PROCEDURE FOR TEACHME EVALUATION

1. Psychologist: Requests participant completes consent, demographics, computer and robot experiences and 10-item personality Inventory (TIPI) questionnaires
2. Experimenter: Explains purpose of experiment, nature of sensorized house and capabilities of robot
3. Experimenter: Explains sets of behaviors to be taught to the robot by the participant
4. Experimenter: Shows participant how to use the TeachMe GUI by jointly creating the “answering the doorbell” behavior and then testing it
5. Experimenter: Ask participant to choose 1 behavior from each of the three sets of behaviors (3 in total) and use the TeachMe GUI to teach the robot each behavior. Once taught, test and modify if necessary
6. Psychologist: Requests participant to complete System Usability Scale and Ad-Hoc questionnaires
7. Psychologist: Participant invited to ask questions and comment on the experience

TABLE III
SYSTEM USABILITY SCALE QUESTIONNAIRE

<i>Modified Brooke’s Usability Scale (5 point Likert scale)</i>
I think that I would like to use the robot teaching system like this often.
I found using the robot teaching system too complex.
I thought the robot teaching system was easy to use.
I thought the robot teaching system was easy to use
I think that I would need the support of a technical person who is always nearby to be able to use this robot teaching system.
I found the various functions in the robot teaching system were well integrated.
I thought there was too much inconsistency in the robot teaching system.
I would imagine that most people would very quickly learn to use the robot teaching system.
I found the robot teaching system very cumbersome to use.
I felt very confident using the robot teaching system.
I needed to learn a lot of things before I could get going with the robot teaching system .

Participants were then asked to choose one behavior from each set of behaviors and use the teaching GUI to teach the robot these behaviors.

During the teaching process, the experimenter stayed with the participant and helped when asked. The post-experimental questionnaire asked them to indicate whether they thought they could continue to use the teaching system without the help of the experimenter. The participant’s use of the teaching system was also videotaped for later analysis.

Having taught the robot a new behavior, the experimenter then invited the participant to test it. If the behavior operated successfully then the participant moved on to teaching the next behavior in the subsequent set. Alternatively, they could modify the existing behavior and retest. Having taught all three behaviors (one from each set), the experimenter retired to another room and the psychologist returned and asked the participant to fill in a post-evaluation questionnaire based on Brooke’s usability scale [46] that had been adapted to the HRI domain from its typical form in HCI (see Table III). A subsequent question-

TABLE IV
COMPUTER USAGE IN THE SAMPLE

Activity	Yes	No
Work or Study	18	2
Socialising	19	1
Recreation	8	12
Programming	0	20

TABLE V
PERSONALITY IN THE SAMPLE

	Mean	SD
Extraversion	4.38	1.48
Agreeable	5.35	1.14
Conscientious	5.83	1.15
Emotional Stability	4.85	1.36
Openness	5.17	1.10

naire (see Table VII) was also completed, which focused on the usefulness of the robot and teaching system specifically.

After completion of the questionnaires, the participant was invited to ask questions. All of the participants were very interested to know how the house and robot worked.

B. Results of the “Teach Me” Evaluation

1) *Demographics*: There were 20 participants in the study: 16 females and four males. The mean age was 44 years, with a median age of 49 years. The age range was from 20 to 67 years. The computer usage of the participants (see Table IV) suggests that the majority of participants used computers for work/studies as well as for social reasons. There was a split in respect to using computers for recreational reasons, such as games. None of the participants programmed computers. The mean number of hours spent on computers was 35 h (SE = 2.98) with a median number of hours of 33. Only one of the participants had had any experience with robots. Table V shows the responses to the TIPI in the sample.

2) *Responses to the “Teach Me” SUS*: The mean participant response to the System Usability Scale regarding the teaching interface was 79.75 (SE = 2.29), and the median response was 76.25. These scores were significantly higher than the “neutral score” of 68 ($t(19) = 5.12, p < 0.001$).

While considering the relationship between the usability scores to the “neutral” score the collaborative carer/primary user usage and training scenarios intended for the “TeachMe/Show Me” system were different from the more industrial settings where the SUS is more commonly applied. As such, the score should be taken as representative of the experienced usability within the interaction context itself rather than merely a representation of the interface [47].

A multiple regression analysis was conducted in order to investigate demographic predictors of SUS responses to this task. After removing nonsignificant predictors, the final model had an adjusted r^2 of 0.28, and predicted SUS scores significantly ($F(2, 17) = 4.70, p = 0.02$). The model is described in

TABLE VI
PREDICTORS OF SUS SCORES

Predictor	β	SE	t(19)	p
Intercept	0.00	0.00	0.00	1.00
Age	-0.49	0.20	-2.48	0.02
Conscientiousness	-0.40	0.20	-2.23	0.04

TABLE VII
FREQUENCIES OF RESPONSES TO THE “TEACH ME” AD-HOC LIKERT ITEMS

Do you think it is useful teach a robot?				
Very Useful	Useful	Neither	Not Useful	Not at all
18	2	0	0	0
Do you think that you would be able to teach the robot?				
Def. Yes	Yes	Neither	No	Def. No
10	10	0	0	0
Would you be willing to teach the robot for someone else e.g. if you were a relative or carer of the other person?				
Def. Yes	Yes	Neither	No	Def. No
14	6	0	0	0
Do you think that robot should already have been completely setup by someone else?				
Def. Yes	Yes	Neither	No	Def. No
1	3	4	11	1
Do you think that the robot should be able to carry out standard tasks but it would be useful to be able to customize it?				
Def. Yes	Yes	Neither	No	Def. No
13	7	0	0	0
Is it useful knowing what the robot can already do?				
Def. Yes	Yes	Neither	No	Def. No
12	8	0	0	0
How would you feel about having a robot suggesting that you take more exercise?				
Very Conf.	Comfortable	Neutral	Unconf.	Very Unconf.
9	8	2	1	0
How would you feel about having a robot suggesting that you play a game together e.g. a video game or chess/draughts?				
Very Conf.	Comfortable	Neutral	Unconf.	Very Unconf.
6	11	2	1	0
How would you feel about having a robot warning you that there was a problem in the house e.g. fridge left open or hot/cold taps running or TV left on?				
Very Conf.	Comfortable	Neutral	Unconf.	Very Unconf.
18	2	0	0	0
How would you feel about having a robot informing someone else that there was a problem in the house e.g. by texting them, if the problem had not been resolved?				
Very Conf.	Comfortable	Neutral	Unconf.	Very Unconf.
12	5	2	1	0

Table VI and suggests that both higher age and higher scores on the Conscientiousness personality trait were associated with lower scores on the SUS for this task.

3) *Responses to the “Teach Me” Ad-Hoc Questions:* Participant responses to the ad-hoc Likert items can be found in Table VII. All participant responded “*Very Useful*” or “*Useful*” when asked if they thought it useful to teach a robot. In addition, all participants answered “*Definitely Yes*” or “*Yes*” when asked if they thought that they would be able to teach the robot, if they would do so for a relative, and that they would find it useful to customize the tasks of a robot beyond a set of standard tasks. The participants did not, however, agree as strongly on whether or not the robot should be completely set up by someone else, with a wider range of responses from the participants.

TABLE VIII
EXPERIMENTAL PROCEDURE FOR SHOWME EVALUATION

1. Experimenter: Explains purpose of experiment
2. Experimenter: Show participant how to use the ShowMe GUI by jointly creating one of the activities in Table IX and then testing to see if it worked
3. Experimenter: Ask participant to choose corresponding task in Table X and teach robot to use this activity. Once taught, test and modify if necessary
4. Psychologist: Requests participant to complete System Usability Scale and Ad-Hoc questionnaires
5. Psychologist: Participant invited to ask questions and comment on the experience

Participants also responded that they were overall “*Very Comfortable*” or “*Comfortable*” with a robot informing them that there was a problem in their house, and 17 out of the 20 participants answered that they were at least “*Comfortable*” with the robot informing a third party about an unresolved problem, but there was less agreement regarding having a robot suggest that they play a game or exercise.

As these were ordinal Likert items, exploratory Spearman’s correlations were carried out.

For wanting the robot already set up, there was a correlation approaching significant between this and the *Emotional Stability* personality trait ($\rho(20) = 0.40, p = 0.08$) indicating that participants with higher scores along this dimension were less likely to want the robot fully set up by someone else. There was also a trend approaching significance for this item and *Age* ($\rho(20) = -0.37, p = 0.10$), in which older participants were more likely to want the robot already set up.

There were no significant relationships between comfort with the robot suggesting that one take more exercise and the demographic measures.

There was a significant relationship between *Age* and Comfort regarding the robot contacting a third party in case of a problem ($\rho(20) = -0.53, p = 0.02$), where older participants were more comfortable with this.

4) *Teaching Behaviors—“Teach Me”—Summary of Results:* Participants found the interface easy to use. Moreover, all participants indicated that they felt able to use a system like this to teach the robot, and willing to use such a system to set-up behaviors for an elderly relative or person in their care.

In terms of individual differences, there are some salient relationships. The relationship between *Age* and SUS scores are not unexpected. The older members of the sample found the system more difficult to use than the younger participants. Related to this is the impact of age on the ad-hoc item regarding wanting the robot to be already set up by someone else. Here, older participants were more likely to want the robot being fully set up than younger participants.

Taken together, the current stage of this teaching system may be better suited for use by carers and relatives of elderly people to set up the robot’s behaviors for them.

The relationship between items covering the possibility of the robot contacting third parties in case of problems, and *Age* is also interesting (and we envisaged that this would be a key item that may be taught to the robot). While one explanation

TABLE IX
SET OF ACTIVITIES USED FOR THE “SHOW ME” EVALUATION

Create an activity called ‘ Watching TV .’
Create an activity called ‘ Relaxing on the Sofa .’
Create an activity called ‘ Preparing a hot drink .’
Create an activity called ‘ Preparing a Ready Meal ’ using the microwave
Create an activity called ‘ Kitchen Activities ’ which is active when ‘Preparing a hot drink’ or ‘Preparing a ready meal’ or when any other kitchen activity is being carried out

TABLE X
SET OF BEHAVIORS TAUGHT AS PART OF THE “SHOW ME” EVALUATION

Teach the robot that if it is 7.30 and you are ‘ Watching TV ’ then remind you that your favourite program is about to start.
If you have been ‘ Relaxing on the sofa ’ for more than 30 min make the robot come to the sofa and tell you to ‘move about a bit’
If there are ‘ Kitchen Activities ’ make the robot come to the kitchen and offer to carry the drink or meal to the dining table

for this result may be that older participants were closer to having to consider these scenarios in their own lives than their younger counterparts, a more likely explanation may be that the older portion of the sample were more likely to have had more experiences with caring for elderly parents or other relatives and so might identify more strongly with the third party that is to be contacted. Some of the informal responses from participants during the debrief of the study did reference such experiences.

C. Learning System—“Show Me”

A preliminary evaluation of the learning system involved three persons, all female aged 58 to 66, who were “informal” carers. They typically looked after an elderly relative. All had previously been exposed to the “Teach Me” system.

1) *Procedure*: The experimenter explained the purpose of the study and ensured that they understood the instructions. The experimenter then chose one of the activities in Table IX and explained to the participant how to use the “Show Me” GUI to allow the robot to learn about this activity—typically by actually carrying out that activity whilst the “start showing me” button was active. They could then test whether this activity was recognized by pressing the ‘test’ button, repeating the activity and ensuring that the recognition bar turned green (i.e., over 50% probable). If lower than 50% the activity was repeated.

The participant was then asked to choose one of the other activities shown in Table IX and use the “Show Me” GUI to allow the robot to learn about the activity. They then tested this activity to ensure that the robot correctly identified it.

Having successfully tested that the robot had learned about this activity, the participant was then asked to choose the corresponding teaching task in Table X. For example, if the robot had learned about “watching TV,” then the behavior involving “watching TV” would be chosen. The participant then taught the robot the chosen behavior and subsequently tested that it worked. For example, for “watching TV,” that the robot would approach the participant (who was now sitting on the sofa watch-

TABLE XI
FREQUENCIES OF RESPONSES TO THE “TEACH ME” AD-HOC LIKERT ITEMS

Do you think it is useful to teach robot activities?				
V. Useful	Useful	Neither	Not Useful	Not at all
3	0	0	0	0
Do you think that <i>you</i> would be able to teach robot activities?				
Def. Yes	Yes	Neither	No	Def. No
2	1	0	0	0
Would you be willing to teach activities for someone else carer of the other person?				
Def. Yes	Yes	Neither	No	Def. No
2	1	0	0	0
Should activities already have been setup by someone else?				
Def. Yes	Yes	Neither	No	Def. No
0	0	0	3	0

TABLE XII
COMMENTS MADE ON THE “SHOW ME” INTERFACE

I think it is a great idea to personalize the robot for an individual’s needs. But also think this can be used alongside prepared repetitive tasks. I think also very important for the robot to learn activities rather than or as well as one off tasks. When teaching activities need to show robot in simple exaggerated steps so that it does not confuse activities
Not completely set up but a range of everyday types of activities which can be personalized

ing TV) and inform them about an upcoming TV program. Following this, the participant was asked to complete the System Usability questionnaire and completed two additional questionnaires on ad-hoc usability and provide, if they wished, an overall comment on the system. A summary is shown in Table VIII.

2) *Results of the “Show Me” Evaluation*: The SUS scores for the “Show Me” interface ranged from 67.5 to 80. The mean score was 75.83 and the median score was 80. This was larger than the expected average of 68. Ad-hoc likert item results are shown in Table XI and some user comments are shown in Table XII.

Clearly, such a small sample may only be indicative, however the results from the SUS suggested that participants found the interface relatively easy to use. The three participants all found the “Show Me” feature useful, and felt confident in their ability to use a feature like this to teach a robot about their own activities or to use on behalf of someone else. They also felt that this should not be something that was already set up prior to use.

V. CONCLUSION

We have described a robot personalization system designed to be used by persons operating in assistive environments in smart homes, typically carers, relatives or the elderly person themselves. The teaching component exploits sets of standard templates in order to generate robot behaviors. This approach avoids the complexity of robot behavior generation for a large set of tasks which we believe would be required by such persons, clearly however more complex tasks would still need technical personnel involvement. The teaching interface was evaluated with end users and indicated that participants considered that

such a system would be both useful and useable by them for aiding persons to stay in their homes for longer periods. We have also described and presented a limited evaluation of a user driven activity learning system which allows the robot and smart home to recognize user activities. This activity recognition system compliments the teaching system by allowing a higher level of semantic behavior creation to be achieved.

The results of both of these studies indicate that such facilities would be readily accepted for use by carers, relatives and the elderly themselves. However, with increasing age, the willingness to learn new ways to operate, by personalizing a robot’s behaviors, decreases.

In these studies the robot was operating primarily as a cognitive prosthetic. However, a question that could be asked is “why use a robot?” and not simply another device such as a mobile phone? We would argue that the use of a robot differs in a number of ways to that of a mobile phone. First, the robot will find the person to inform them (a mobile phone may be somewhere else and ignored). Second, there is some evidence [48]–[51] that the robot, by having a physical presence, is perceived as more authoritative, i.e., a person is more likely to follow a robot’s instructions or suggestions rather than, say, a phone.

The exploration, via the “Show Me” system, of creating higher level semantics, is we believe a novel and promising way to ease the teaching burden. For example, being able to instruct a robot by using everyday terms, such as “when it’s time for bed do...” or “if I’m making dinner do” We have only partially explored such opportunities and issues which surround “showing” a robot typical activities and this work is at an early stage. A number of improvements and enhancements to such facilities would be to use both inductive and predictive mechanisms to increase the reliability of the robot recognizing user activities. Prediction algorithms already exist which use past sensory data to predict possible next actions [52]. A further extension of this work would be to use those predictions to then predict again—effectively creating a predictive forward model for the robot. This forward model then being subject to the inductive algorithm, which would now use both historical and predicted sensor vectors to make a decision on user activity.

This area of research also presents some ongoing design challenges that are currently being pursued from two largely distinct viewpoints. The first viewpoint focuses on people-centered initiatives and improving acceptance by tackling HRI issues by giving control on personalization and product customization features. The second viewpoint studies technologically driven initiatives by building impersonal systems that are able to autonomously adapt their operations to fit changing requirements, but ignore HRI. In order to inform the development of a new generation of smart robotic spaces, solutions to the combination of these different research strands is, we believe, a fundamental requirement.

Finally, we have demonstrated in this work that personalization of an autonomous robot is possible in a domestic environment.

REFERENCES

- [1] B. Przywara, “Projecting future health care expenditure at European level: drivers, methodology and main results,” in *European Economy*, European Commission, Economic and Financial Affairs, 2010.
- [2] Eurostats. (2013). Population projections. [Online]. Available: http://epp.eurostat.ec.europa.eu/statistics_explained
- [3] Welsh Social Services Improvement Agency. (2012, Nov. 23). Demonstrating improvement through reablement. [Online]. Available: <http://www.ssiacymru.org.uk/reablement>
- [4] M. K. Lee, J. Forlizzi, S. Kiesler, P. Rybski, J. Antanitis, and S. Savetsila, “Personalization in HRI: A longitudinal field experiment,” in *Proc. Int. Conf. Human-Robot Interaction*, 2012, pp. 319–326.
- [5] U. Reiser, C. Connette, J. Fischer, J. Kubacki, A. Bubeck, F. Weisshardt, T. Jacobs, C. Parlitz, M. Hagele, and A. Verl, “Care-o-bot creating a product vision for service robot applications by integrating design and technology,” in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2009, pp. 1992–1998.
- [6] G. Pilkington, *Homecare Re-Ablement*. London, U.K.: Dept. Health, 2008.
- [7] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, “A survey of robot learning from demonstration,” *Robot. Auton. Syst.*, vol. 57, no. 5, pp. 469–483, May 2009.
- [8] S. Chernova and A. L. Thomaz, *Robot Learning from Human Teachers. Synthesis Lectures on Artificial Intelligence and Machine Learning*. San Rafael, CA, USA: Morgan & Claypool, 2014.
- [9] M. Nicolescu and M. J. Mataric, “Learning and interacting in human-robot domains,” *IEEE Trans. Syst., Man, Cybern. A, Syst. Humans*, vol. 31, no. 5, pp. 419–430, Sep. 2001.
- [10] R. E. Fikes and N. J. Nilsson, “Strips: A new approach to the application of theorem proving to problem solving,” *Artif. Intell.*, vol. 2, pp. 189–208, 1971.
- [11] J. Saunders, N. Otero, and C. L. Nehaniv, “Issues in human/robot task structuring and teaching,” in *Proc. IEEE Int. Symp. Robot Human Interactive Commun.*, 2007, pp. 708–713.
- [12] C. Sammut, S. Hurst, D. Kedzier, and D. Michie, “Learning to fly,” in *Proc. 9th Int. Conf. Mach. Learning*, San Mateo, CA, USA: Morgan Kaufmann, 1992, pp. 385–393.
- [13] C. Sammut, S. Hurst, D. Kedzier, and D. Michie, “Learning to fly,” in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds., Cambridge, MA, USA: MIT Press, 2002, pp. 171–190.
- [14] J. R. Quinlan, *C4.5: Programs for Machine Learning*. San Mateo, CA, USA: Morgan Kaufmann, 1993.
- [15] M. van Lent and J. E. Laird, “Learning procedural knowledge through observation,” in *Proc. 1st Int. Conf. Knowl. Capture*, Victoria, BC, Canada, 2001, pp. 179–186.
- [16] C. L. Nehaniv and K. Dautenhahn, “The correspondence problem,” in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds., Cambridge, MA, USA: MIT Press, 2002, pp. 41–61.
- [17] A. Levas and M. Selfridge, “A user-friendly high-level robot teaching system,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 1984, vol. 1, pp. 413–416.
- [18] M. Kaiser, V. Klingspor, J. del R. Millán, M. Accame, F. Wallner, and R. Dillmann, “Using machine learning techniques in real-world mobile robots,” *IEEE Expert*, vol. 10, no. 2, pp. 37–45, Apr. 1995.
- [19] M. Kaiser and R. Dillman, “Building elementary robot skills from human demonstration,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 1996, vol. 3, pp. 2700–2705.
- [20] H. Friedrich and R. Dillman, “Robot programming based on a single demonstration,” presented at the 3rd Eur. Workshop Learning Robots, Crete, Greece, 1995.
- [21] R. Dillmann, “Teaching and learning of robot tasks via observation of human performance,” *Robot. Auton. Syst.*, vol. 47, pp. 109–116, 2004.
- [22] A. L. Thomaz and C. Breazeal, “Teachable robots: Understanding human teaching behavior to build more effective robot learners,” *Artif. Intell.*, vol. 172, pp. 716–737, Apr. 2008.
- [23] Y. Kuniyoshi, M. Inaba, and H. Inoue, “Learning by watching: Extracting reusable task knowledge from visual observations of human performance,” *IEEE Trans. Robot. Autom.*, vol. 10, no. 6, pp. 799–822, Nov. 1994.
- [24] K. Ikeuchi and T. Suehiro, “Towards an assembly plan from observation,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 1992, pp. 2171–2177.
- [25] M. Johnson and Y. Demiris, “Hierarchies of coupled inverse and forward models for abstraction in robot action planning, recognition and imitation,” in *Proc. Int. Symp. Animals Artifacts*, Apr. 2005, pp. 69–76.

- [26] K. S. Narendra and J. Balakrishnan, "Adaptive control using multiple models," *IEEE Trans. Autom. Control*, vol. 42, no. 2, pp. 171–187, Feb. 1997.
- [27] D. M. Wolpert and M. Kawato, "Multiple paired forward and inverse models for motor control," *Neural Netw.*, vol. 11, pp. 1317–1329, 1998.
- [28] T. Gu, Z. Wu, X. Tao, H. Pung, and J. Lu, "epsICAR: An emerging patterns based approach to sequential, interleaved and concurrent activity recognition," in *Proc. Int. Conf. Pervasive Comput. Commun.*, 2009, pp. 1–9.
- [29] N. C. Krishnan and D. J. Cook, "Activity recognition on streaming sensor data," *Pervasive Mobile Comput.*, vol. 10, Part B, no. 0, pp. 138–154, 2014.
- [30] J. Ye, S. Dobson, and S. McKeever, "Situation identification techniques in pervasive computing: A review," *Pervasive Mobile Comput.*, vol. 8, no. 1, pp. 36–66, 2009.
- [31] J. Saunders, C. L. Nehaniv, and K. Dautenhahn, "Teaching robots by moulding behavior and scaffolding the environment," in *Proc. Conf. Human-Robot Interaction*, Mar. 2006, pp. 118–125.
- [32] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source robot operating system," in *Proc. ICRA Workshop Open Source Softw.*, 2009.
- [33] N. Hu, G. Englebienne, and B. J. A. Kröse, "Bayesian fusion of ceiling mounted camera and laser range finder on a mobile robot for people detection and localization," in *Proc. IROS Workshop: Human Behavior Understanding*, 2012, pp. 41–51.
- [34] N. J. Nilsson, "Teleo-reactive programs for agent control," *J. Artif. Intell. Res.*, vol. 1, pp. 139–158, 1994.
- [35] J. Saunders, M. Salem, and K. Dautenhahn, "Temporal issues in teaching robot behaviours in a knowledge-based sensorised home," presented at the 2nd Int. Workshop Adaptive Robot. Ecologies, Dublin, Ireland, 2013.
- [36] D. Nau, Y. Cao, A. Lotem, and H. Muñoz-Avila, "SHOP: Simple Hierarchical Ordered Planner," in *Proc. 16th Int. Joint Conf. Artif. Intell.*, 1999, pp. 968–973.
- [37] R. Hartanto, *A Hybrid Deliberative Layer for Robotic Agents: Fusing DL Reasoning with HTN Planning in Autonomous Robots*. Berlin, Germany: Springer, 2011.
- [38] D. Off and J. Zhang, "Multimodal integration processes in plan-based service robot control," *Tsinghua Sci. Technol.*, vol. 16, no. 1, pp. 1–6, 2011.
- [39] J. Saunders, N. Burke, K. L. Koay, and K. Dautenhahn, "A user friendly robot architecture for re-ablement and co-learning in a sensorised homes," *Assistive Technology: From Research to Practice*, P. Encarnacao, Ed. Amsterdam: IOS Press, 2013, pp. 49–58.
- [40] N. J. Nilsson, "Teleo-reactive programs and the triple-tower architecture," *Electron. Trans. Artif. Intell.*, vol. 5, pp. 99–110, 2001.
- [41] C. Chen, B. Das, and D. Cook, "A data mining framework for activity recognition in smart environments," in *Proc. Int. Conf. Intell. Environ.*, 2010, pp. 80–83.
- [42] J. Lester, T. Choudhury, N. Kern, G. Borriello, and B. Hannaford, "A hybrid discriminative/generative approach for modelling human activities," in *Proc. 19th Int. Joint Conf. Artificial Intell.*, 2005, pp. 766–772.
- [43] J. Saunders, C. L. Nehaniv, and K. Dautenhahn, "Using self-imitation to direct learning," in *Proc. IEEE Ro-man 2006 15th IEEE Int. Workshop Robot Human Interactive Commun.*, Sep. 2006, pp. 244–250.
- [44] I. Duque, K. Dautenhahn, K. L. Koay, I. Willcock, and B. Christianson, "Knowledge-driven user activity recognition for a smart house. development and validation of a generic and low-cost, resource-efficient system," in *Proc. 6th Int. Conf. Adv. Comput.-Human Interactions*, 2013, pp. 141–146.
- [45] S. D. Gosling, P. J. Rentfrow, and W. B. Swann Jr, "A very brief measure of the big-five personality domains," *J. Res. Personality*, vol. 37, no. 6, pp. 504–528, 2003.
- [46] J. Brooke, "SUS-A quick and dirty usability scale," *Usability Evaluation Ind.*, vol. 189, p. 194, 1996.
- [47] J. Brooke, "SUS: A retrospective," *J. Usability Stud.*, vol. 8, no. 2, pp. 29–40, 2013.
- [48] J. Wainer, D. J. Feil-Seifer, D. A. Shell, and M. J. Mataric, "Embodiment and human-robot interaction: A task based perspective," in *Proc. IEEE Int. Conf. Human-Robot Interaction*, 2007, pp. 872–877.
- [49] H. Kose-Bagci, E. Ferrari, K. Dautenhahn, D. S. Syrdal, and C. L. Nehaniv, "Effects of embodiment and gestures on social interaction in drumming games with a humanoid robot," *J. Adv. Robot.*, vol. 23, pp. 1951–1996, 2009.
- [50] W. Bainbridge, J. Hart, E. Kim, and B. Scassellati, "The benefits of interactions with physically present robots over video-displayed agents," *Int. J. Soc. Robot.*, vol. 3, no. 1, pp. 41–52, 2011.
- [51] M. Salem, G. Lakatos, F. Amirabdollahian, and K. Dautenhahn, "Would you trust a (faulty) robot? effects of error, task type and personality on human-robot cooperation and trust," in *Proc. IEEE Int. Conf. Human-Robot Interaction*, pp. 141–148, 2015.
- [52] K. Gopalratnam and D. J. Cook, "Active lezi: An incremental parsing algorithm for sequential prediction," in *Proc. 16th Int. Florida Artif. Intell. Res. Soc. Conf.*, 2003, pp. 38–42.



Joe Saunders received the Ph.D. degree from the University of Hertfordshire, Hertfordshire, U.K., in 2007.

He is currently a Senior Research Fellow with the Adaptive Systems Research Group, University of Hertfordshire. He has carried out research on robot adaptation and learning at both physical and linguistic levels.



Dag Sverre Syrdal received the B.Sc. degree in psychology from Queen's University of Belfast, Belfast, U.K., and the M.Sc. degree in research methods from the University of Hertfordshire, Hertfordshire, U.K.

His research interests include individual differences in human interaction, evaluation methodologies, and social aspects of human-robot interaction.



Kheng Lee Koay received the Ph.D. degree from the University of Plymouth, Plymouth, U.K., in 2003.

He is currently a Senior Research Fellow with the Adaptive Systems Research Group, University of Hertfordshire, Hertfordshire, U.K. His research interests include mobile robotics, robotic home companions, and human-robot interaction.



Nathan Burke received the B.Sc. degree in computer science from Texas State University, San Marcos, TX, USA, and the M.Sc. degree in artificial intelligence from the University of Hertfordshire, Hertfordshire, U.K. He is currently working toward the Ph.D. degree in artificial intelligence with the University of Hertfordshire.

His research interests include machine learning, robot teaching methods, and assistive robotics.



Kerstin Dautenhahn (SM'13) received the Ph.D. degree from the University of Bielefeld, Bielefeld, Germany, in 1993.

She is currently a full Professor with the School of Computer Science, University of Hertfordshire, Hertfordshire, U.K., where she coordinates the Adaptive Systems Research Group.