

# Towards Reliable Autonomous Robotic Assistants Through Formal Verification: A Case Study

Matt Webster, Clare Dixon, Michael Fisher, Maha Salem,  
Joe Saunders, Kheng Lee Koay, Kerstin Dautenhahn (*Senior Member, IEEE*) and Joan Saez-Pons

**Abstract**—It is essential for robots working in close proximity to people to be both safe and trustworthy. We present a case study on formal verification for a high-level planner/scheduler for the Care-O-bot<sup>®</sup>, an autonomous personal robotic assistant. We describe how a model of the Care-O-bot<sup>®</sup> and its environment was developed using Brahms, a multi-agent workflow language. Formal verification was then carried out by automatically translating this model to the input language of an existing model checker. Four sample properties based on system requirements were verified. We then refined the environment model three times to increase its accuracy and the persuasiveness of the formal verification results. The first refinement uses a user activity log based on real-life experiments, but is deterministic. The second refinement uses the activities from the user activity log non-deterministically. The third refinement uses “conjoined activities” based on an observation that many user activities can overlap. The four samples properties were verified for each refinement of the environment model. Finally, we discuss the approach of environment model refinement with respect to this case study.

**Index Terms**—human-robot teams, formal verification, autonomous systems, model checking, robotics.

## I. INTRODUCTION

Robotic assistants are likely to be used for a variety of applications including personal healthcare, exploration within remote environments, and manufacturing. These robots will operate in close proximity to their human operators and therefore must be safe and trustworthy in their operations. One of the aims of the EPSRC-funded *Trustworthy Robotic Assistants* (TRA) project<sup>1</sup> is to develop tools and techniques for the verification and validation of robotic assistants. The TRA project uses three different methodologies for this: formal verification, simulation-based testing and end-user validation. In this paper we present a case study on the application of formal verification to the Care-O-bot<sup>®</sup>: an autonomous robotic



Fig. 1. The Care-O-bot<sup>®</sup> Robotic Assistant operating in the University of Hertfordshire’s Robot House.

assistant deployed at the University of Hertfordshire’s Robot House (Fig. 1).

Verification is the process of assessing whether a system meets its requirements; formal verification is the application of formal (i.e., mathematical) methods to the verification of systems. The formal verification approach used in this paper is based on model checking [1], in which a model of a program or process is constructed. This model is typically non-deterministic, so that each ‘run’ (or simulation) of the model can be different from the last. A program called a model checker exhaustively analyses all possible executions of the model in order to establish that some property, usually derived from system requirements, holds. Therefore it is possible, for example, to use a model checker to formally verify that in every execution of a given program, the program will always reach a desirable situation. In other words, we can formally verify that a given requirement holds.

The model checker used for this case study, SPIN [2], has been publicly available since 1991 and has been used for the formal verification of a wide variety of systems, including flood control barriers, telecommunications switches and several space missions [3]. SPIN, which stands for Simple PROMELA Interpreter, verifies programs and processes written in PROMELA, the Process Meta-Language. Rather than writing in PROMELA directly, we utilise an intelligent agent modelling language and simulation environment called Brahms [4] to develop models of the Robot House and Care-O-bot<sup>®</sup>. Brahms can be used to develop detailed models of systems with multiple interacting agents and has been used to model human-robot teams [5] and complex workflows [4] for space exploration. We translate the Brahms models automati-

Matt Webster is a Postdoctoral Research Associate in the Department of Computer Science, University of Liverpool, Liverpool, L69 3BX, UK. Email: matt@liverpool.ac.uk.

Clare Dixon is a Senior Lecturer in the Department of Computer Science, University of Liverpool.

Michael Fisher is a Professor in the Department of Computer Science, University of Liverpool.

Maha Salem is a Postdoctoral Research Fellow in the School of Computer Science, University of Hertfordshire, Hertfordshire, AL10 9AB, UK.

Joe Saunders is a Senior Research Fellow in the School of Computer Science, University of Hertfordshire.

Kheng Lee Koay is a Senior Research Fellow in the School of Computer Science, University of Hertfordshire.

Kerstin Dautenhahn is a Professor in the School of Computer Science, University of Hertfordshire.

Joan Saez-Pons is a Postdoctoral Research Fellow in the School of Computer Science, University of Hertfordshire.

<sup>1</sup><http://www.robosafe.org/>

cally into PROMELA using the *BrahmsToPromela* translator software [6] based on a formal semantics for Brahms [7]. The autonomous control systems used in the Robot House were written at a similar level of abstraction to constructs in the Brahms modelling language. Therefore encoding these autonomous systems using Brahms reduced modelling errors due to changes in abstraction level.

Brahms is a rich language for specifying the behaviour of multi-agent systems, and includes features such as inheritance, geography, message passing and probabilities. Whilst some of the features were necessary in modelling the Robot House (e.g., message passing), others were useful but not strictly necessary (e.g., inheritance). It would have been possible to use a lower level language to model the Robot House environment directly, e.g., PROMELA, and therefore avoid the additional step needed to translate from Brahms to PROMELA. However, as Brahms workframes closely resembled the if-then rules used in the autonomous control systems within the Robot House, modelling these rules using Brahms could be achieved with limited effort. In addition, the *BrahmsToPromela* translator had already been developed, and therefore the quickest way to obtain a PROMELA model from the autonomous control system code was to first encode the if-then rules as Brahms workframes, and then translate into PROMELA automatically using *BrahmsToPromela*.

In the remainder of this section we examine the Care-O-bot<sup>®</sup> and the Robot House in more detail, and compare our approach to related work. In Section II we describe the way in which a model of the Care-O-bot<sup>®</sup>'s high-level planner/scheduler for autonomous decision-making system was developed using Brahms. Then, in Section III, we show how that model was formally verified using the SPIN model checker, and in Section IV we improve the environment models used by the model checker. Finally, in Section V, we provide a concluding discussion.

#### A. The Robot House and Care-O-bot<sup>®</sup>

The University of Hertfordshire's Robot House is a suburban three-bedroom house near Hatfield, UK (see Figs. 1 and 2). In addition to house furnishings and décor, the Robot House is equipped with more than 50 sensors which provide real-time episodic information on the state of the house and the individuals occupying it. These sensors range across electrical (e.g., refrigerator door open/closed sensor), furniture (e.g., cupboard drawers open), services (e.g., detect when toilet flush is being used) and pressure (e.g., chair sensors to detect when someone is seated) devices [8], [9].

The Robot House hosts a number of different robots that are used to conduct Human-Robot Interaction (HRI) experiments (e.g., [10]). One of these robots is the commercially-available Care-O-bot<sup>®</sup> robot manufactured by Fraunhofer IPA [11]. It has been specifically developed as a mobile robotic assistant to support people in domestic environments, and is based on the concept of a 'robot butler' [12]. The Care-O-bot<sup>®</sup> robot is equipped with a 7 degrees-of-freedom manipulator arm extended with a three-finger gripper and also comprises an articulated torso, stereo sensors serving as 'eyes', LED lights



Fig. 2. A plan view of the ground floor of the University of Hertfordshire Robot House. Numbered boxes show the locations of sensors.

and a tray. Accordingly, the robot's sensors include its current location, the state of the arm, torso, eyes and tray. By means of a text-to-speech synthesising module, the robot is also capable of expressing given text as audio output.

The robot's software is based on the Robot Operating System (ROS) and a number of ROS packages (e.g., drivers, navigation and simulation software) are available online<sup>2</sup>. For example, to navigate to any designated location within the house, the robot uses the ROS navigation package<sup>2</sup> in combination with its laser range-finders to perform self-localisation, map updating, path planning, and obstacle avoidance in real-time while navigating along the planned route.

High-level commands are sent to the robot via the ROS *script server* mechanism which are then interpreted into low-level commands by the robot's software. For example, these high-level commands can take the form 'raise tray', 'move to location  $x$ ', 'grab object on tray', 'say hello', etc.

The Care-O-bot<sup>®</sup>'s high-level decision making is determined by rules (stored in a MySQL database) of the form:

```
Guard
==
RobotAction*
```

Here, *Guard* is a sequence of propositional statements that are either true or false, linked by Boolean AND (&) and OR (|) operations, while *RobotAction\** is a sequence of actions which the Care-O-bot<sup>®</sup> will perform only if the *Guard* is true. In practice, the *Guard* is implemented as a set of SQL queries and the *RobotActions* are implemented through the ROS-based *cob\_script\_server* package, which provides a simple interface to operate Care-O-bot<sup>®</sup>. For example, take the following rule which lowers the Care-O-bot<sup>®</sup>'s tray:

```
SELECT * FROM Sensors WHERE sensorId=500
```

<sup>2</sup><http://wiki.ros.org/care-o-bot>, <http://wiki.ros.org/navigation>

```

AND value = 1      &
SELECT * FROM Sensors WHERE sensorId=504
AND value = 1
==
light,0,yellow
tray,0,down,,wait
light,0,white,,wait
cond,0,500,0
cond,0,501,1

```

The SQL table called `Sensors` stores the values of all sensors in the Robot House. One row at most is stored in the Sensor table for each `sensorId / value` pair, so that the SELECT queries above will return exactly one value if the proposition they represent is true, and false otherwise. (The word ‘sensors’ is used loosely and includes variables which describe the Robot House’s knowledge about the state of the house, e.g., the ‘trayRaised’ variable.)

In the rule above, the guard checks whether variable 500 (‘trayRaised’) and 504 (‘trayEmpty’) are set to 1 (‘true’) or not by performing the SQL SELECT queries. If the guard is true, the Care-O-bot<sup>®</sup> will change the light colour to yellow (indicating that the robot is in motion), set the tray to the lowered position, wait for completion, set the light to white (indicating that the robot has stopped moving), and wait for completion. Finally the variables 500 (‘trayRaised’) and 501 (‘trayLowered’) are set to ‘false’ and ‘true’ respectively. Note that twice the robot `waits` for short periods of time (around one second). The `wait` command is used to prevent the robot’s actions from executing at the same time, e.g., to prevent the light being set to white (indicating to the user that the robot has stopped moving) before the robot actually stops moving.

The Care-O-bot<sup>®</sup>’s rule database is composed of multiple rules for determining a variety of autonomous behaviours, including checking and answering the front doorbell and reminding a person to take medication. The full Robot House rule database includes a set of 31 default rules and can be obtained from the EU ACCOMPANY project’s Git repository<sup>2</sup>.

### B. Related Work

Stocker et al. [6] describe the development of the *Brahms-ToPromela* software tool, which is utilised in this work to automatically translate a model of the Care-O-bot<sup>®</sup> written in Brahms into a PROMELA specification which can then be formally verified using the SPIN model checker. The authors examine an assisted living scenario similar to the Robot House system tackled in this case study. However, the work here expands on the work of Stocker et al. by modelling a real-life robotic system and scenario where the rules are directly derived from actual code used in practice. Additionally, this work has used user activity logs from real-world experiments within the Robot House in order to increase the verisimilitude of the person agent within the model.

Saunders et al. [8] and Duque et al. [9] described the University of Hertfordshire Care-O-bot<sup>®</sup> and Robot House systems which are used as a basis for this work. More information on the development of these systems can be found on the ACCOMPANY project website<sup>3</sup>.

Formal verification has been used before for robotic systems. For example, Mohammed et al. [13] used hybrid automata and hybrid statecharts for formal modelling and verification of multi-robot systems; Cowley and Taylor [14] used dependent type theory and linear logic for the formal verification of assembly robots; and Kouskoulas et al. [15] formally verified control algorithms for surgical robots. However, very little work has been conducted to formally verify the safety and trustworthiness of robotic home companions. This is where our work aims to complement existing research in the area of formal verification of autonomous and robotic systems.

## II. MODELLING THE CARE-O-BOT<sup>®</sup> USING BRAHMS

The 31 default rules are similar in structure to common constructs within the Brahms multi-agent workflow programming language. The first step in modelling was to convert the full set of Care-O-bot<sup>®</sup> rules into a more convenient *if-then* rule representation. For example, the above rule was rewritten as:

```

IF tray_is_raised AND tray_is_empty
THEN set_light(yellow)
    move_tray(lowered_position)
    wait()
    set_light(white)
    wait()
    set(tray_is_raised,false)
    set(tray_is_lowered,true)

```

Once translated into this format, these rules could then be straightforwardly translated into Brahms. A key concept in Brahms is the ‘workframe’, which specifies a sequence of things to be done when a given condition holds. The Robot House rules were translated into Brahms workframes within the Care-O-bot<sup>®</sup> agent, with the IF a THEN b rules translated into the when a do { b } construct in Brahms. For example, the rule above was translated into a Brahms workframe called `wf_lowerTray`:

```

workframe wf_lowerTray {
  repeat: true;
  priority: 10;
  when(knownval(current.trayIsRaised = true) and
    knownval(current.trayIsEmpty = true))
  do{
    conclude((current.lightColour =
      current.colourYellow));
    lowerTray();
    wait();
    conclude((current.lightColour =
      current.colourWhite));
    wait();
    conclude((current.trayIsRaised = false));
    conclude((current.trayIsLowered = true));
  }
}

```

This workframe is set to ‘repeat’, meaning that it can be used more than once by the agent. Multiple workframes can be eligible for execution by the Brahms interpreter at the same time, so the `priority` sets the importance of the workframe relative to other workframes (with larger numbers being more important). The `when a do { b }` construct states that when some conditions are true, in this case `trayIsRaised` and `trayIsEmpty`, the agent should do the actions that follow. In the action list, the `conclude()` construct is used to determine when beliefs should be updated within the Brahms agent. The

<sup>2</sup><https://github.com/uH-adapsys/accompany>

<sup>3</sup><http://accompanyproject.eu>

`lowerTray()` and `wait()` statements are programmer-defined primitive actions whose function is to denote that something has happened and has a certain duration, e.g.:

```
primitive_activity wait(){ max_duration: 1; }
```

In general, the Robot House rules were translated into Brahms workframes on a one-to-one basis. However, in some cases it was necessary to use more than one Brahms workframe for a rule. This generally happened when a rule contained interactions with the user via the GUI. For instance, when the person sits down and watches television (detected via sensors in the sofa seats and the television power outlet) the Care-O-bot<sup>®</sup> approaches the person and asks whether he or she would like to watch the television together. The person has three options which are presented by the Robot House via a GUI on a tablet computer: to tell the Care-O-bot<sup>®</sup> to watch television, to return to its charging station, or to continue with its current task. This behaviour is modelled using a Brahms workframe within the Care-O-bot<sup>®</sup> agent, in which these options are communicated to the person using the `announceQueryToUser_ThreeOptions()` activity:

```
workframe wf_watchTV {
  repeat: true;
  priority: 10;
  when(knownval(robotHouse.sofaSeatOccupied =
    true) and
    knownval(robotHouse.televisionWattage > 10)
    and ...)
  do{
    conclude((current.queryUserOption1 =
      current.activityWatchTV));
    conclude((current.queryUserOption2 =
      current.activityReturnHome));
    conclude((current.queryUserOption3 =
      current.activityContinue));
    conclude((current.userQueried = true));
    conclude((current.queryUser_ThreeOptions = true));
    announceQueryToUser_ThreeOptions();
    conclude((current.AskedToWatchTV = true));
  }
}
```

The ‘person’ agent (simulating simple human behaviour) then selects a response and sends it back to the Care-O-bot<sup>®</sup> model. In this example, the person agent was set to always choose to watch TV with the Care-O-bot<sup>®</sup> *when asked* by the Care-O-bot<sup>®</sup>. The following Care-O-bot<sup>®</sup> workframe is executed when the person decides to watch TV:

```
workframe wf_optionSelectedWatchTV {
  repeat: true;
  priority: 10;
  when(knownval(Person.userRespondedToQuery = true)
    and
    knownval(Person.queryResponse =
      current.activityWatchTV))
  do{
    conclude((Person.userRespondedToQuery = false));
    conclude((Person.guiSetWatchTV = true));
    conclude((current.userQueried = false));
  }
}
```

Here, the guard uses the Care-O-bot<sup>®</sup>’s belief about the person agent’s belief about whether the person has responded to the query. (While the belief is referred to as `Person.userRespondedToQuery`, it concerns the Care-O-bot<sup>®</sup>

agent’s belief about the person agent’s belief as this workframe is part of the Care-O-bot<sup>®</sup> agent. For more information on belief handling in Brahms see [16].). If this belief is set to ‘true’, and the person has responded to the query, then the workframe will execute. Note that later in the workframe the Care-O-bot<sup>®</sup> agent sets this belief to false so that the next time the person agent responds to a query, this workframe will be able to execute correctly.

### A. Modelling a Scenario

After translating the Robot House rules into Brahms it was necessary to set up a model of the Robot House environment, or *scenario*. The scenario determines the range of possibilities within the Robot House environment consisting primarily of the Care-O-bot<sup>®</sup>, the person being assisted by Care-O-bot<sup>®</sup>, and the Robot House. For example, the scenario consists of a model of the person and the Robot House, where each is defined as an agent within Brahms. Another agent, the ‘Campanile Clock’, measures the passage of time in the model and keeps the other agents updated with the current time.

Information within the Brahms model is stored as agent beliefs, which can be public or private. Agents can reason about their own beliefs as well as the beliefs of other agents. For example, the Campanile Clock agent maintains the current time within the model as a public belief, and periodically announces the time to the other agents. Therefore the Campanile Clock’s `time` variable can be used by all agents to refer to the current time within the model. (This is useful when defining properties for formal verification.) Other information in the model is held by the relevant agent. For example, the Robot House agent maintains sensor data for Robot House, the Care-O-bot<sup>®</sup> agent maintains information about the location and state of the Care-O-bot<sup>®</sup>, and the Person agent maintains the location and state of the person occupying the Robot House.

The layout of the Robot House is encoded using a *geography* within Brahms, providing an hierarchical description of the different places that the agent can occupy. For example:

```
areadef House extends BaseAreaDef { }
areadef Room extends House { }
areadef areaOfInterest extends Room { }
area LivingRoom instanceof Room partof House { }
area LivingRoomTV instanceof areaOfInterest
  partof LivingRoom { }
```

This Brahms geography states that there is a type of area called `House`. `House` is an extension of a built-in Brahms class called `BaseAreaDef`. Another area type, `Room`, extends the `House` class, and `areaOfInterest` extends `Room`. The `LivingRoom` is an instance of the `Room` class, and is part of the `House`, and the `LivingRoomTV` is an instance of `areaOfInterest` and is part of the `LivingRoom`. Note that the `House` area definition is different from the Robot House agent described above; the former describes the layout of the robot house, while the latter encapsulates the data from the Robot House’s sensors together with functionality for communicating sensor state with other agents.

Our scenario lasts from noon to 9pm. At any given point in the day the person may choose to sit and watch TV, move into the living room area, or move into the kitchen (e.g., to

prepare food to eat), or may choose to send the Care-O-bot<sup>®</sup> into the kitchen or the living room. At 5pm the person needs to take medication.

The person can act non-deterministically, that is, behave in a manner which is unpredictable within the model of the scenario. This non-determinism is implemented in Brahms as a set of five workframes within the person agent, all of which will fire at a given point. Each workframe has a priority. The highest priority workframe is executed, and a belief is modified within the agent (using the ‘conclude’ keyword in Brahms). This belief is modified with a level of certainty (known as the belief-certainty) which states that the belief will be updated with a given probability. If the belief is updated, this information is communicated to the Care-O-bot<sup>®</sup> agent or the Robot House agent (depending on the workframe) which causes these agents to know that the person has done something, e.g., sent the Robot to the kitchen via the GUI, or that the person has moved into the living room. If the belief is not updated, then the next workframe fires. It is possible for none of the five workframes to update a belief within the person agent, and this special case models the ability of the person to do nothing.

Based on this simple scenario we can establish a number of high-level requirements for the Care-O-bot<sup>®</sup>. For example, the Care-O-bot<sup>®</sup> should remind the person to take medication at 5pm. Another requirement is that the Care-O-bot<sup>®</sup> should go into the living room if it is told to go into the living room by the person. In the next section, we formalise these kinds of requirements using temporal logic and verify them using the SPIN model checker.

### III. FORMAL VERIFICATION OF BRAHMS MODELS USING *BrahmsToPromela* AND SPIN

Brahms refers to a multi-agent workflow specification language, as well as a software package consisting of an agent simulation toolkit and an integrated development environment. The Brahms software does not come with formal verification tools built-in; for formal verification we used the *BrahmsToPromela* translator [6]. *BrahmsToPromela* allows models written using a subset of Brahms to be automatically translated into PROMELA, the process meta-language used by the SPIN model checker. Once translated, SPIN can be used for the automatic formal verification of the Brahms model with respect to particular requirement. In our case, we formalise these requirements using linear temporal logic (LTL), which allows the formalisation of concepts relating to time, e.g., ‘now and at all points in the future’ (via the  $\Box$  operator), ‘now or at some point in the future’ ( $\Diamond$ ) and ‘in the next state’ ( $\bigcirc$ ) [17]. This enables formalisation of safety requirements (something bad never happens,  $\Box \neg \text{bad}$ ), liveness properties (e.g. something good eventually happens,  $\Diamond \text{good}$ ) and fairness properties (e.g., if one thing occurs infinitely often so does another, e.g.,  $\Box \Diamond \text{send} \implies \Box \Diamond \text{receive}$ ).

Using *BrahmsToPromela* extends SPIN’s property specification language with a *belief* operator, ‘B’. This is parameterised by the agent that holds the belief — so  $B_{\text{Care-O-bot } x}$  means that the Care-O-bot<sup>®</sup> believes  $x$  is true. Beliefs in Brahms

are translated into boolean variable arrays in PROMELA. For example,  $B_{\text{Care-O-bot } x}$  is modelled by “`bool CareOBot_x[n]`”, where  $n$  is the number of agents and objects within the Brahms simulation. Beliefs are stored as arrays as Brahms allows agents to (dis)believe other agents’ beliefs, so it is necessary to use an array to store whether or not each agent believes a particular agent’s belief. Using this framework we model  $B_{\text{Care-O-bot}}(B_{\text{Person } x})$  (i.e., whether the Care-O-bot<sup>®</sup> believes that the Person believes  $x$ ) as the Boolean array index “`Person_x[cob]`” where `cob` is a constant that refers to the Care-O-bot. Similarly,  $B_{\text{Person } x}$  is modelled as “`Person_x[person]`”.

All of the properties in this paper use beliefs rather than facts about the environment. Beliefs may or may not hold, i.e., beliefs can be incorrect. However, in this system we assume complete sensor accuracy so that if  $x$  is true, then  $B_a x$  is true for all agents  $a$ .

To explore possibilities, the following sample requirements were translated and their formalised properties verified using SPIN for the Brahms model.

- 1) It is always the case that if the Care-O-bot<sup>®</sup> believes that the person has told it to move into the kitchen, then it will eventually move into the kitchen.

$$\Box \left[ B_{\text{Care-O-bot}}(B_{\text{Person}} \text{guiGoToKitchen}) \implies \Diamond B_{\text{Care-O-bot}}(\text{location} = \text{Kitchen}) \right]$$

- 2) It is always the case that if the Care-O-bot<sup>®</sup> believes that the person has told it to move to the sofa in the living room, then it will eventually move there.

$$\Box \left[ B_{\text{Care-O-bot}}(B_{\text{Person}} \text{guiGoToSofa}) \implies \Diamond B_{\text{Care-O-bot}}(\text{location} = \text{LivingRoomSofa}) \right]$$

- 3) It is always the case that if the Robot House believes that the sofa seat has been occupied for at least one hour<sup>4</sup>, and if the Robot House believes that the television power consumption is higher than 10 watts, and if the Care-O-bot<sup>®</sup> believes that it has not yet asked the person if he or she wants to watch television, then eventually the Care-O-bot<sup>®</sup> will move to the living room sofa and ask the person if he or she wants to watch the television.

$$\Box \left[ B_{\text{RobotHouse}} \text{SofaOccupied1Hour} \wedge B_{\text{RobotHouse}} \text{tvPower} > 10 \wedge B_{\text{Care-O-bot}} \neg \text{askedToWatchTV} \implies \Diamond \left( B_{\text{Care-O-bot}} \text{location} = \text{LivingRoomSofa} \wedge B_{\text{Care-O-bot}} \text{askedToWatchTV} \right) \right]$$

- 4) It is always the case that if the time is 5pm, then the Care-O-bot<sup>®</sup> will believe that the medicine is due.

$$\Box \left[ B_{\text{Campanile\_Clock}} \text{time} = 5\text{pm} \implies \Diamond B_{\text{Care-O-bot}} \text{medicineDue} \right]$$

The first two requirements are derived from a higher-level requirement that, in general, the Care-O-bot<sup>®</sup> should follow instructions given to it by the person. The third property is important for maintaining the social activity of the person within the Robot House, who is temporarily under the care of the Care-O-bot<sup>®</sup>. The fourth property is derived from a higher-level requirement that the Care-O-bot<sup>®</sup> should remind the person to take medication at the correct time.

<sup>4</sup>This ensures that the Care-O-bot<sup>®</sup> will only ask the person if he or she wants to watch television once every hour at most.

Table I summarizes the verification results<sup>5</sup>. These results were obtained using an eight-core Intel® Core™ i7-3720QM CPU (2.60GHz) laptop with 16 GB of memory running Ubuntu Linux 12.04 LTS. In each case the same PROMELA model was used; any difference in the number of states or time taken is due to the complexity of the property being verified and the resulting automaton used by the model checker. For requirements 1, 2 and 4 the resources used were almost identical, and this is to be expected as these properties were similar in structure and produced similar automata. Property 3 produced a slightly more complex automaton requiring more resources to verify.

TABLE I  
FORMAL VERIFICATION OF FOUR PROPERTIES FOR THE ORIGINAL MODEL.

Prop.	States	Depth	Memory (MB)	Time (s)
1	302, 160	74, 819	399	9.9
2	302, 160	74, 819	399	9.7
3	576, 317	81, 341	410	18.0
4	302, 160	74, 819	399	9.7

#### IV. IMPROVING THE ENVIRONMENT MODEL

In the previous section we were able to formally verify the high level decision making system within the Care-O-bot®. This provides a degree of assurance that the Care-O-bot® will satisfy the requirements/properties against which it has been verified. However, this degree of assurance is dependent on the quality of the model of the Care-O-bot®'s environment. The environment model determines the behaviour of all agents beyond the Care-O-bot®, including any people within the Robot House, the Robot House and the sensors within the Robot House. In the case of the Care-O-bot®, which has a deterministic control system, the environment model is the sole source of non-determinism within the model and generates the state space which is then model checked.

However, the environment model used in the Brahms model is simple. It describes a scenario that lasts from 12pm to 9pm. In each hour the person agent can choose from one of six possible options: sit down and watch TV, move into the living room area, move into the kitchen, send the Care-O-bot® into the kitchen, send the Care-O-bot® to the living room, or do nothing. At 5pm the person will need to take medication, and is reminded to do so by the Care-O-bot®. Which one of the six possible options is chosen by the person is non-deterministic, and during simulation of the model the choice made is random.

In a real-world Robot House scenario there are many more things that could happen than in our environment model. The person may go to bed, for example, or choose to leave the house. The doorbell could ring, or another person could enter the house.

In order to generate a more interesting environment model we used the activity recognition system in [9] to collect data from an actual person living in the Robot House for a period of

four days. The person's behaviour was monitored by sensors throughout the Robot House (see Fig. 2). There are a total of 59 sensors in the Robot House which allow the house to be monitored in a variety of ways. A real-time energy monitoring system can detect the activation and de-activation of electrical devices like refrigerators and kettles. A different sensor network is able to monitor the movement of people though the Robot House using reed sensors, temperature sensors and pressure mats. Via these sensors it is possible to detect when someone sits down on a chair, or opens the bedroom door, for example [9].

One person occupied the house for approximately four days and eleven hours. Combinations of sensor outputs were used to determine the current activity of the person within the Robot House during this time. For instance, if the 'bed contact' sensor is activated, then it was assumed that the `In_Bed` activity was underway. In other words, the person must be in bed. Activities can also be related to other activities. For example, if the activity `In_Bed` is active, and the activity `Lamp_on_Bedroom` has been deactivated for at least 5 seconds, then the activity `Sleeping_Bedroom` is inferred. If the person is in bed, and has just turned the light off, then the person must be sleeping in the bed.

The user activity log generated automatically in the Robot House contained 569 different entries, each being a tuple containing the activity name, whether the activity was beginning (i.e., being activated) or finishing (i.e., being deactivated) and the date/time. For example, the following state that the activity `In_Bed` began at 11:03pm on 19th May 2013 and finished at 5:25am on 20th May 2013:

```
In_Bed 1 2013-05-19 23:03:50
In_Bed 0 2013-05-20 05:25:03
```

The entries covered a range of activities, including `In_Bathroom`, `Toileting`, `Preparing_Hot_Drink`, `Actively_Watching_TV` and `Water_Sink_ON`.

Three different approaches to modelling the person were taken in light of the user activity log. They are as follows.

##### A. Deterministic Environment Model

In the deterministic environment model the entries in the user activity log were converted directly into Brahms workframes without abstraction, so that if an event occurred at a particular time within the Robot House, then the same event occurs within the Brahms model at exactly the same time. For example, take the following entries in the activity log:

```
Preparing_Hot_Drink 1 2013-05-20 09:07:06
Preparing_Food 1 2013-05-20 09:07:06
```

Both activities start at the same time, 9:07:06 am. These were translated into the following Brahms workframe within the person agent:

```
workframe wf_36196 {
  repeat: true;
  priority: 100;
  when(knownval(Campanile_Clock.time = 36196 )
    and knownval(current.wf36196 = false))
  do {
    conclude((current.PreparingHotDrink = true));
    conclude((current.PreparingFood = true));
```

<sup>5</sup>The computational resources used are less than those reported in our earlier paper [18] due to the use of hash compression in the SPIN model checker. For consistency, hash compression was used throughout this paper.






Fig. 3. Part of the UH Robot House user activity log for a six-hour period beginning at approximately 4am on Thursday 23rd May 2013.

```

conclude((current.wf36196 = true));
wait();
}
}

```

This workframe was set to repeat (although repetition would not happen due to the way the workframe is defined) with priority 100, the same as all other workframes generated from the activity log. The `when` statement says that when the campanile clock time is set to 36196<sup>6</sup>, to set the person agent’s beliefs about preparing a hot drink and preparing food to true, and to set the belief `wf36193` to true. A precondition of this `when` statement is that `wf36196` is false; setting it to `true` prevents this workframe from executing again in the future.

A custom Java application automated the process to translate a comma-delimited input file containing the complete activity log into a set of Brahms workframes of the form above. Brahms agent belief declarations, such as `PreparingFood`, were also generated by the Java application. Since all beliefs were describing the activity of the person within the Robot House, these belief declarations were included in the Person agent’s specification in Brahms.

The resulting Brahms file was 316 kB in size and contained over five hundred workframes in the person agent — including one for each user activity log entry. This is larger than the earlier Brahms model (which was  $\sim 107$  kB), and Brahms was able to execute the model in 10 seconds on the same laptop used for model checking. However, translating this Brahms model into SPIN resulted in a model too large to be model checked. There were too many `mtype` elements (symbolic names for constant values) within the PROMELA model. (SPIN allows a maximum of 255 distinct names within an `mtype` declaration.) The `mtype` elements are used by the

*BrahmsToPromela* translator to keep track of agent states within the PROMELA model. One strategy was to see whether it was possible to verify a fragment of the user activity log instead.

A fragment corresponding to the first 21 hours of user activity was the largest fragment that could be formally verified. The Robot House model incorporating the deterministic environment model was formally verified with respect to the four properties given in Section III. Table II summarizes the time and memory usage.

TABLE II  
FORMAL VERIFICATION OF FOUR PROPERTIES FOR DETERMINISTIC ENVIRONMENT MODEL

Prop.	States	Depth	Memory (MB)	Time (s)
1	114,501	229,001	1,463	8.4
2	45,272	90,543	674	3.3
3	109,329	218,655	1,394	8.1
4	51,009	102,017	742	3.6

The model checker was able to formally verify these properties much more quickly (a minimum of 3.3s compared to 9.7s), and with fewer states (a minimum of 45,272 compared to 302,160). This is to be expected, as the model is deterministic, resulting in a much smaller state space.

From a model checking perspective, the deterministic environment model can be seen as a single run (or simulation) of a non-deterministic model. However, this model is still interesting from a formal verification perspective as it allows validation of the robot’s autonomous control system as modelled within Brahms and translated using *BrahmsToPromela*. For example, we can observe the behaviour of the real-world Care-O-bot<sup>®</sup> over the 21 hour period contained in the user activity log, and check whether properties 1–4 actually held in the real-world. If they did not, this would indicate that

<sup>6</sup>The clock time of 36196 was based on 9:07:06 am being 36,196 seconds after the first entry in the activity log.

there was a problem with the Brahms model and that re-design would be necessary. Furthermore, use of a deterministic model within a model checker allows us to verify a formally-defined property, and therefore provide a more rigorous result than simply executing a simulation of the Brahms model and checking the output trace to see whether the requirement corresponding to the formal property has held.

### B. Extended Non-Deterministic Environment Model

One of the strongest aspects of model checking is to automatically explore the state space of a non-deterministic model. Within such a model, there are points at which a number of different things could happen. For example, a person may sit down, watch TV, or prepare a meal. These non-deterministic choices produce a branching state space which typically grows exponentially with the number of such choices. One job of a model checker, therefore, is to check all the possible paths through this state space in order to ensure that a particular property holds at all points during the execution of the model. For example, in Section III we showed that the Care-O-bot<sup>®</sup> will always remind the person in the house to take their medication at 5pm. By using a model checker, we know that this is always true in every possible run through the model, and in every possible state, with respect to the environment model used.

In order to improve the environment model in Section II-A another non-deterministic model was developed. This new model operates in a similar manner: there are a number of possible actions that the person can take, and the choice of which action to take is non-deterministic. However, this environment model is a significant improvement over the environment model in Section II-A as it contains a higher number of possible actions.

In the environment model from Section II-A, there were only six things the person could choose to do in a given time step. Inspection of the user activity logs revealed that there were many more than six things that the person did and these activities were used as the basis for the new environment model.

A fragment of the user activity log can be seen in Fig. 3. Each of the activities shown on the left hand side of the figure were modelled using Brahms. There were a total of 26 different user activities from which the person agent could choose (22 activities shown in Fig. 3, and four user activities from the environment model in Section II that were not duplicated in the user activity log).

The multi-agent version of the Robot House scenario, including the extended non-deterministic environment model described above, was model checked with respect to the four properties given in Section III. No errors were found. The time and memory requirements are summarized in Table III.

Increased non-determinism in the environment model has resulted in the use of more computational resources during model checking. For example, the number of states required has roughly tripled compared to the results in Table I. This is to be expected as there are more activities available to the person agent.

TABLE III  
FORMAL VERIFICATION OF FOUR PROPERTIES FOR EXTENDED  
NON-DETERMINISTIC ENVIRONMENT MODEL


Prop.	States	Depth	Memory (MB)	Time (s)
1	1, 122, 836	113, 519	1, 223	107
2	1, 122, 836	113, 519	1, 223	104
3	1, 300, 161	113, 519	1, 118	111
4	1, 122, 836	113, 519	1, 223	104

### C. Non-deterministic Conjoined Activity Environment Model

Examination of the user activity log showed that a number of activities overlap, and were not mutually-exclusive. (A number of overlapping activities can also be seen in Fig. 3.) Consider the following:

```
In_Bathroom 1 2013-05-20 05:25:46
Toileting 1 2013-05-20 05:26:01
Toileting 0 2013-05-20 05:26:56
In_Bathroom 0 2013-05-20 05:27:05
```

Here the `In_Bathroom` activity begins, and a few seconds later the `Toileting` activity begins and then ends. Shortly after the `In_Bathroom` activity ends. Clearly these two activities have overlapped. We could describe what has happened with a simple state diagram:



Here `none` denotes a state in which no activities are happening. In the next state, `In_Bathroom` is active, and in the next state, `InBathroom_Toileting` is happening. Next `In_Bathroom` occurs again (as `Toileting` has now finished) and finally we return to the `none` state.

The non-deterministic model in Section IV-B is therefore less accurate than we thought, as the user can select from 26 mutually exclusive activities, whereas in reality these activities are not mutually exclusive and may overlap. In order to improve the accuracy of the environment model we therefore needed to account for overlapping activities. This was done by ‘conjoining’ the activities in the way we did with the `In_Bathroom` and `Toileting` activities above: these overlapping activities can be converted into the mutually exclusive states `none`, `In_Bathroom` and `InBathroom_Toileting`. In other words, by conjoining the names of the currently-occurring activities we can describe them as mutually-exclusive states of the form we had in the example in Section II-A.

For the conversion, a Java application was implemented. The application reads in an activity log from a comma-delimited input file. This is then converted into a data structure *D* which maps timestamps to sets of activities which are active at that time. For example, for the activity log excerpt above, *D* would be:



```

{ 2013-05-2005:25:46 ⇨ {InBathroom},
  2013-05-2005:26:01 ⇨ {InBathroom,Toileting},
  2013-05-2005:26:56 ⇨ {InBathroom},
  2013-05-2005:27:05 ⇨ {} }

```

Each set in the range of  $D$  is converted into a string denoting a conjoined activity, so that  $\{InBathroom,Toileting\}$  becomes `InBathroom_Toileting`. The conversion is performed deterministically so that each set maps to a single conjoined activity string, with no possibility of having, say, `InBathroom_Toileting` distinct from `Toileting_InBathroom`. This set of strings is then used to output Brahms workframes describing those conjoined activities. For example, the `SittingDiningArea`, `HavingMeal` and `MedicineTaken` activities all overlap and so were converted to the conjoined activity `SittingDiningArea_HavingMeal_MedicineTaken`, in turn translated into the following Brahms workframes:

```

workframe wf_SittingDiningArea_HavingMeal_Medicine-
Taken {
  repeat: true;
  priority: 25;
  when(knownval(Campanile_Clock.time < 5) and
    knownval(current.doneSittingDiningArea_-
      HavingMeal_MedicineTaken = false) and
    knownval(current.consideredSitting-
      DiningArea_HavingMeal_MedicineTaken
      = false))
  do {
    conclude((current.consideredSittingDiningArea_-
      HavingMeal_MedicineTaken = true));
    conclude((current.goalDoSittingDiningArea_-
      HavingMeal_MedicineTaken = true),bc:20);
    wait();
  }
}

workframe wf_doSittingDiningArea_HavingMeal_-
MedicineTaken {
  repeat: true;
  priority: 25;
  when(knownval(current.goalDoSittingDiningArea_-
    HavingMeal_MedicineTaken = true) and
    knownval(current.doneSittingDiningArea_-
      HavingMeal_MedicineTaken = false))
  do {
    conclude((current.doneSittingDiningArea_-
      HavingMeal_MedicineTaken = true));
    conclude((current.SittingDiningArea = true));
    conclude((current.HavingMeal = true));
    conclude((current.MedicineTaken = true));
    wait();
  }
}

```

The first workframe says that if the time is before 5pm, and the conjoined activity `SittingDiningArea_HavingMeal_MedicineTaken` has not been done, and if this conjoined activity has not yet been ‘considered’, then consider it by making a belief update with certainty 20% (`bc:20`). Therefore, 20% of the time this update is successful, which will cause the next workframe to execute and update the beliefs of the person agent to reflect that the activities `SittingDiningArea`, `HavingMeal` and `MedicineTaken` are now taking place. 80% of the time this belief update is not successful, which does not trigger the second workframe. In this case, another workframe

corresponding to a distinct conjoined activity may execute. Ultimately, it is possible that none of the conjoined activities happens, in which case the person agent is doing nothing for that time step.

Using the ActivityLog Java application a total of 133 conjoined activities were found. Each of which was converted into Brahms workframes as described above.

The multi-agent system of the Robot House scenario, including the non-deterministic conjoined activity environment model, was model checked with respect to the four properties given in Section III. No errors were found. The time and memory requirements are presented in Table IV.

TABLE IV  
FORMAL VERIFICATION OF FOUR PROPERTIES FOR NON-DETERMINISTIC  
CONJOINED ACTIVITY ENVIRONMENT MODEL

Prop.	States	Depth	Memory (MB)	Time (s)
1	4, 470, 619	337, 437	5, 354	872
2	4, 248, 201	373, 391	5, 938	871
3	4, 587, 998	384, 157	6, 190	859
4	4, 430, 294	372, 549	5, 926	848

This environment model contained a larger number of activities for the person to be engaged in than in the non-deterministic environment model used previously (there are now 133 different conjoined activities compared to 26 different activities previously). It can be seen that the increased non-determinism in the environment model has increased the computational resources required by the model checker. For example, the number of states used is around four times as many as for the extended non-deterministic environment model, and around 14 times as many as for the environment model used in Section III.

## V. DISCUSSION

In this case study, we formally verified an autonomous decision making planner/scheduler system for the Robot House assisted living environment and the Care-O-bot<sup>®</sup> robotic assistant. This was done by converting the Robot House planner/scheduler rules into Brahms: a workflow language for defining the behaviour of multiple agents. These rules matched closely the Brahms style. Brahms was also used to model the Robot House environment, including the Care-O-bot<sup>®</sup>, the Robot House and a resident. The Brahms model was then translated into the PROMELA language using the *BrahmsToPromela* tool [6]. Once in PROMELA, the Brahms workflows (and, consequently, the Care-O-bot<sup>®</sup>’s decision making system) were formally verified. Four properties were formally verified, demonstrating that this process can be used for verification of autonomous decision making systems for robotic systems.

The simplistic non-deterministic environment model used in Section III was enhanced in three different ways in order to increase its fidelity. First it was enhanced to a deterministic model that covered a larger set of user activities. Then, it was modified to include a total of 26 different activities for the person agent to choose from (20 activities more than the initial non-deterministic environment model). Finally, it was modified to allow for the overlapping of the person agent’s

activities, resulting in 133 different conjoined activities. In all cases the model checker showed that there were no errors found, therefore providing assurance that the Care-O-bot®'s high level decision-making system meets requirements.

The person agent in the environment model is the chief source of non-determinism, as is the case with the resident in the real-life Robot House. Extending the environment model in three different ways (to give a total of four different environment models) means that the unpredictability of the person agent has been modelled in four different ways. Each of the four models provide a degree of assurance that the Care-O-bot® is safe for use by a single person in the Robot House, and together they provide an even greater degree of assurance. Obviously a model is an abstraction of the real world and will never be able to fully represent all aspects of the actual situation. However, by modelling human behaviour in as many ways as possible, we gain a greater confidence in the system being formally verified, and a greater level of trust in the autonomous robotic assistant.

The aim of this case study is to provide a detailed example of how formal verification can be applied to an existing personal robotic system. We show how formal verification can be used to provide assurance that the robotic system will behave correctly with respect to a small subset of the robot's requirements. Whilst the size of this subset is small, it is sufficient as a proof-of-concept.

The formal verification techniques used were effectively 'off-the-shelf' software components, consisting of Brahms, a multi-agent simulation framework, *BrahmsToPromela*, an automatic translator from Brahms to PROMELA, and SPIN, a well-known model checker for the PROMELA language. The use of modular components expedites the modelling and formal verification process, and minimises errors that may be introduced by developing new formal verification methods.

#### ACKNOWLEDGMENT

This work was supported by EPSRC grants EP/K006193 and EP/K006509.

#### REFERENCES

- [1] E. M. Clarke, O. Grumberg, and D. A. Peled, *Model Checking*. USA: MIT Press, 1999.
- [2] G. J. Holzmann, *The SPIN Model Checker: Primer and Reference Manual*. USA: Addison-Wesley, 2003.
- [3] G. Holzmann, "Inspiring applications of Spin," <http://spinroot.com/spin/success.html>, 2013, accessed 2013-10-10.
- [4] M. Sierhuis and W. J. Clancey, "Modeling and Simulating Work Practice: A Method for Work Systems Design," *IEEE Intelligent Systems*, vol. 17, no. 5, pp. 32–41, 2002.
- [5] W. J. Clancey, M. Sierhuis, C. Kaskiris, and R. van Hoof, "Brahms mobile agents: Architecture and field tests," in *Human-Robot Interaction: Papers from the AAAI Fall Symposium*, 2002, ISBN 978-1-57735-174-0.
- [6] R. Stocker, L. A. Dennis, C. Dixon, and M. Fisher, "Verification of Brahms Human-Robot Teamwork Models," in *Proc. 13th European Conf. on Logics in Artificial Intelligence*, ser. LNCS, vol. 7519. Springer, 2012, pp. 385–397.
- [7] R. Stocker, M. Sierhuis, L. A. Dennis, C. Dixon, and M. Fisher, "A Formal Semantics for Brahms," in *Proc. 12th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA)*, ser. LNCS, vol. 6814. Springer, 2011, pp. 259–274.
- [8] J. Saunders, N. Burke, K. L. Koay, and K. Dautenhahn, "A User Friendly Robot Architecture for Re-ablement and Co-learning in A Sensorised Home," in *Proc. 12th European AAATE Conference*, 2013.

- [9] 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. Sixth International Conference on Advances in Computer-Human Interactions (ACHI)*, Nice, France, 2013, pp. 141–146.
- [10] D. Syrdal, K. Dautenhahn, K. L. Koay, M. Walters, and W. Ho, "Sharing Spaces, Sharing Lives – The Impact of Robot Mobility on User Perception of a Home Companion Robot," in *Proc. Fifth International Conference on Social Robotics (ICSR)*, ser. LNAI. Springer, 2013, vol. 8239.
- [11] U. Reiser, C. Connette, J. Fischer, J. Kubacki, A. Bubeck, F. Weisshardt, T. Jacobs, C. Parltitz, M. Hägele, and A. Verl, "Care-O-bot® 3: Creating a Product Vision for Service Robot Applications by Integrating Design and Technology," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE Press, 2009, pp. 1992–1998. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1733023.1733071>
- [12] U. Reiser, T. Jacobs, G. Arbeiter, C. Parltitz, and K. Dautenhahn, "Care-O-bot® 3 — Vision of a Robot Butler," in *Your Virtual Butler — The Making-of*, ser. LNAI, R. Trapp, Ed. Springer, 2013, vol. 7407, pp. 97–116.
- [13] A. Mohammed, U. Furbach, and F. Stolzenburg, "Multi-robot systems: Modeling, specification, and model checking," *Robot Soccer*, pp. 241–265, 2010.
- [14] A. Cowley and C. J. Taylor, "Towards Language-based Verification of Robot Behaviors," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2011, pp. 4776–4782.
- [15] Y. Kouskoulas, D. W. Renshaw, A. Platzer, and P. Kazanzides, "Certifying the Safe Design of a Virtual Fixture Control Algorithm for a Surgical Robot," in *Proc. Hybrid Systems: Computation and Control (HSCC)*, C. Belta and F. Ivancic, Eds. ACM, 2013, pp. 263–272.
- [16] A. Acquisti, W. J. Clancey, R. van Hoof, M. Scott, and M. Sierhuis, *Brahms Tutorial TM01-0002, Version 1.2*, Agent iSolutions, March 2011.
- [17] M. Fisher, *An Introduction to Practical Formal Methods Using Temporal Logic*. Wiley, 2011.
- [18] M. Webster, C. Dixon, M. Fisher, M. Salem, J. Saunders, K. L. Koay, and K. Dautenhahn, "Formal verification of an autonomous personal robotic assistant," in *Formal Verification and Modeling in Human-Machine Systems: Papers from the AAAI Spring Symposium (FVHMS 2014)*, 2014.



**Matt Webster** is a postdoctoral Research Associate in the Department of Computer Science at the University of Liverpool. He is currently working on the Trustworthy Robotic Assistants project funded by EPSRC. He is a member of the Centre for Autonomous Systems Technology and the Logic and Computation research group at Liverpool. His research interests include autonomous systems, formal verification, robotics and unmanned aircraft.



**Clare Dixon** is a Senior Lecturer in the Department of Computer Science at the University of Liverpool. Her research interests include formal specification and verification and temporal and modal theorem-proving techniques. She is a member of the Logic and Computation research group at Liverpool, and she is a member of the Centre for Autonomous Systems Technology, the Institute for Risk and Uncertainty and is deputy director of the Liverpool Verification Laboratory.



**Michael Fisher** is Professor of Computer Science and Director of the Centre for Autonomous Systems Technology (<http://www.liv.ac.uk/cast>) at the University of Liverpool. His research interests involve logical methods in Computer Science and Artificial Intelligence, particularly temporal reasoning, programming languages, practical formal verification, and the development and analysis of autonomous and agent-based systems.



**Joan Saez-Pons** is a Postdoctoral Research Fellow in the Adaptive Systems Research Group at the University of Hertfordshire. His research interests are within the broad area of robotics, with particular interest in autonomous mobile robots, coordination and cooperation of multi-robot systems, human robot interaction, artificial intelligence, social robotics and computer vision.



**Maha Salem** is a Postdoctoral Research Fellow in the Adaptive Systems Research Group at the School of Computer Science, University of Hertfordshire, UK. Her research revolves around social human-robot interaction (HRI) with particular focus on safety aspects for the design of trustworthy robotic assistants. She is further interested in cross-cultural HRI as well as in multimodal interaction and non-verbal expressiveness based on the design and use of communicative robot gesture.



**Joe Saunders** is a Senior Research Fellow in the Adaptive Systems Group of the University of Hertfordshire (UH), UK. He is especially interested in applying theoretical studies to practical real-world environments and has recently been carrying out trials on end-user robot teaching in a care environment as part of the European Union ACCOMPANY project.



**Kheng Lee Koay** is a Senior Research Fellow in the Adaptive Systems Research Group (School of Computer Science, Faculty of Engineering and Information Sciences) at the University of Hertfordshire. He received his B.Sc. degree in Robotics and Automated Systems and Ph.D. degree from the University of Plymouth, U.K. in 1997 and 2003, respectively. His research interests include Mobile Robotics, Human-Robot Interaction, Social Robotics, Home Companion and Agent Migration.



**Kerstin Dautenhahn** is a Professor of Artificial Intelligence in the School of Computer Science at University of Hertfordshire, where she is a coordinator of the Adaptive Systems Research Group. Main areas of her research are Human-Robot Interaction, Social Robotics, Socially Intelligent Agents and Artificial Life.