

# Specification, Stochastic Modeling and Analysis of Interactive Service Robotic Applications

Livia Lestingi<sup>a,\*</sup>, Davide Zerla<sup>a</sup>, Marcello M. Bersani<sup>a</sup>, Matteo Rossi<sup>b</sup>

<sup>a</sup>*Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Via Ponzio 34/5, 20133 Milan, Italy.*

<sup>b</sup>*Dipartimento di Meccanica, Politecnico di Milano, Via Privata Giuseppe La Masa 1, 20156 Milan, Italy.*

---

## Abstract

Assistive robotic systems are quickly becoming a core technology for the service sector as they are understood capable of supporting people in need of assistance in a wide variety of tasks. This step poses a number of ethical and technological questions. The research community is wondering how service robotics can be a step forward in human care and aid, and how robotics applications can be realized in order to put the human role at the forefront. Therefore, there is a growing demand for frameworks supporting robotic application designers in a “human-aware” development process. This paper presents a model-driven framework for analyzing and developing human-robot interactive scenarios in non-industrial settings with significant sources of uncertainty. The framework’s core is a formal model of the agents at play—the humans and the robot—and the robot’s mission, which is then put through verification to estimate the probability of completing the mission. The model captures non-trivial features related to human behavior, specifically the unpredictability of human choices and physiological aspects tied to their state of health. To foster the framework’s accessibility, we present a verification tool-agnostic Domain-Specific Language that allows designers lacking expertise in formal modeling to configure the interactive scenarios in a user-friendly manner. We compare the formal analysis outputs with results obtained by deploying benchmark scenarios in the physical environment with a real mobile robot to assess whether the formal model adheres to reality and whether the verification results are accurate. The entire development pipeline is then tested on several scenarios from the healthcare setting to assess its flexibility and effectiveness in the application design process.

---

## 1. Introduction

Breakthrough technological advancements are shaping the future of the service sector. Innovations brought by the phenomenon known as Industry 4.0, such as IoT, pervasive sensorization, Cloud Computing, and Collaborative Robotics, are now spreading to non-industrial settings with significant projected impacts on our everyday lives. Most importantly, highly sophisticated robotic systems under development today are bound to transform the job market once they become commercially available. The uptake of such solutions poses a number of problems which range from technological challenges to ethical and societal implications. A recent study on the future of employment indeed estimates that specific jobs, such as receptionists, information clerks, healthcare support workers, and personal care aides, will be taken over by robots with probabilities ranging from 60% to 90% [1]. In addition, the presence of robots in healthcare has increased in recent years and shows an accelerating trend [2]. The use and penetration of robotics for human care and aid is evidenced by the presence European calls and projects<sup>1</sup>, technology companies<sup>2</sup>, and market analysis reports [3, 4].

---

\*Corresponding author

*Email addresses:* livia.lestingi@polimi.it (Livia Lestingi), davide.zerla@mail.polimi.it (Davide Zerla), marcellomaria.bersani@polimi.it (Marcello M. Bersani), matteo.rossi@polimi.it (Matteo Rossi)

<sup>1</sup>Examples are the Harmony (<https://harmony-eu.org>) and the EnrichME (<https://cordis.europa.eu/project/id/643691>) projects.

<sup>2</sup>Examples are Kompai Robotics (<https://kompairobotics.com>) and Labrador Systems (<https://labradorsystems.com>).

19 All these initiatives and companies agree that the use of robots in care can increase the quality of services  
20 provided, although robots are not a substitute for humans, but a tool to improve their actions. For instance,  
21 a study by the American Nurses Association [5] showed that robots can support and augment nursing care  
22 delivery, improves nurse productivity, increases time with patients, encourages positive emotional responses.  
23 Despite these evidences, the investigation of the extent of such a technological and societal shift is an ongoing  
24 research. This work attempts to answer the question on the feasibility of such a step by addressing the  
25 analysis from the software engineering standpoint and, in particular, sheds light on the development of  
26 collaborative service robot applications in healthcare.


27 State-of-the-art technologies dealing with sensing, manipulation, and reasoning capabilities make it  
28 feasible for robots to perform complex jobs. Nowadays, a robot may be adequately equipped to sense multiple  
29 aspects of its surroundings, efficiently detect obstacles, grasp and manipulate fragile objects, perform surgery,  
30 and make decisions in delicate situations. However, these skills usually constitute silos of software, whose  
31 integration and reuse are challenging tasks. The EFFIROB project [6], which analyzed the profitability  
32 of developing a new service robot application, has estimated that up to 80% of the total cost comes from  
33 software development and maintenance.

34 More generally, software engineering techniques for robotics are not mature yet to handle the complexity  
35 and changeability of service settings [7]. Service robots operate in unconstrained environments where humans,  
36 who they frequently interact with, constitute a significant source of uncertainty. Decisions made at an early  
37 design stage of the application determine up to 90% of the overall life-cycle costs [8], and considerable sources  
38 of uncertainty can hinder their validity. Therefore, it is of paramount importance to provide designers  
39 with frameworks to develop applications that are simultaneously **reliable** and **flexible** with respect to the  
40 variability of the environment [9]. Frameworks should also limit the gap between the developer’s knowledge  
41 and the prerequisites needed to access them, removing the barriers that are due to the lack of specialized  
42 skills in the developers.

### 43 1.1. Model-Driven Framework

44 Designing robotic applications to be deployed in delicate environments where robots will closely interact  
45 with humans is a challenging task, requiring a strong technical background both in robotics and software  
46 design. Our work contributes to this line of research by proposing a **model-driven framework** to develop  
47 interactive service robot applications. Target *users* of the framework, called hereafter robotic application  
48 designers (or, simply, application designers), are professional figures managing the logistics of service facilities  
49 where robotic applications will be deployed, such as clinical workflow analysts [10]. The framework targets  
50 robotic applications set in known layouts, featuring a wheeled mobile robot and one (or multiple) humans  
51 requesting a service that requires interaction or coordination with the robot. While the geometry of the  
52 layout is known, humans are a source of uncertainty as they may make unpredictable choices and stray from  
53 the plan while interacting with the robot. Applications eligible for analysis come, though not exclusively,  
54 from the healthcare and assisted living settings, where people might be in pain or discomfort. Therefore, the  
55 development process encompasses features of human **physiological** (i.e., physical fatigue) and **behavioral**  
56 aspects, such as the unpredictability of the human decision-making process.

57 Within the scope of the framework, interactions between a human and a robot conform to high-level  
58 “*patterns*” identifying recurring contingencies in assistive applications. Throughout the paper, we use term  
59 “*action*” to indicate an “*atomic behaviour that is executed by any actor in a scene*” [11]. Each interaction  
60 pattern is a sequence of actions (e.g., move until a certain event occurs, stop, wait for the human to be  
61 close). Although there is no standard definition of robotic “*mission*”, with this term we refer to a sequence  
62 of interaction patterns identifying the desired behavior of the robot [12] performed in a specific layout. A  
63 sequence of missions constitutes a Human-Robot Interaction (HRI) “*scenario*”. Hence, in the scope of this  
64 work, we understand a robotic application as the realization of a scenario through real or virtual agents.  
65 The framework exploits formal analysis to provide the robotic application designer with reliable insights  
66 into the outcome of each mission (each analyzed individually) constituting the scenario. Given the initial  
67 configuration of a scenario (e.g., positions of the agents, battery charge), the application designer receives an  
68 estimation of how likely the associated missions are to end in *success* (dually, in *failure*) and the physical  
69 effort each mission imposes on human subjects.



(a) Agents (represented in their starting locations) and POIs of the running example. (b) Layout for the running example, highlighting the two areas, their dimensions ([m]), and the intersection point.

Figure 1: Graphical representation of the example scenario configuration.

*Example.* Fig. 1 shows the setup of a possible scenario. The layout is a T-shaped corridor made up of two rectangular areas (see Fig. 1b): a horizontal one and a perpendicular vertical one, whose intersection is centered in point  $(45.0, 12.5)$ . The corridor features four Points Of Interest (POIs), i.e., significant locations within the layout, also represented in Fig. 1a: the robot’s recharge station (RC), two cupboards containing medical kits (referred to as KIT1 and KIT2), and the door leading to the waiting room (WR). There are four agents in the scenario: two humans (HUM1 and HUM2) and two robots (ROB1 and ROB2). Robot ROB1 is a Tiago<sup>3</sup> with initial battery charge equal to 40% of the total capacity and ROB2 is a Turtlebot3 WafflePi<sup>4</sup> with 90% of the total capacity. HUM1 is a young patient with average walking speed 80cm/s while HUM2 is a healthy elderly doctor with average walking speed 100cm/s. The designer assesses two alternative missions to determine which one is most likely to succeed within a given time bound and with no harm for the humans: the first mission features ROB1 leading HUM1 to the waiting room, then delivering KIT2 to HUM2. The second mission features ROB2 following HUM2 to fetch KIT1, then leading HUM1 to the waiting room.

The framework’s workflow (shown in Fig. 2) is structured into three phases:

- (1) **design-time analysis:** the robotic application designer configures the scenario through a specification language. Starting from the configuration, a formal model of the scenario is automatically generated together with a set of properties. Such properties are subject to verification to estimate quality measures of the scenario (e.g., the probability of success);
- (2) **deployment:** when the design-time results are deemed acceptable, the application designer deploys the scenario either in a physical environment or simulated environment; to enable deployment, the formal model is converted into executable code communicating with the deployment environment through a middleware layer;
- (3) **reconfiguration:** the application designer examines quality metrics of the scenario estimated from data collected during deployment and applies reconfiguration measures, if necessary (e.g., in the first mission of the example, ROB2 might be deployed instead of ROB1 because of the higher charge level).

## 1.2. Contributions

This paper builds upon the results presented in [13, 14, 15, 16] by presenting:

1. A custom **Domain-Specific Language (DSL)** for scenario configuration. Since application designers possibly lack a strong background in software development, formal modeling, or formal analysis, the DSL, which is a lightweight textual notation, provides a friendly interface to the configuration phase.

<sup>3</sup>Technical specifications available at: <https://pal-robotics.com/robots/tiago/>.

<sup>4</sup>Technical specifications available at: <https://emanual.robotis.com/docs/en/platform/turtlebot3/overview/>.

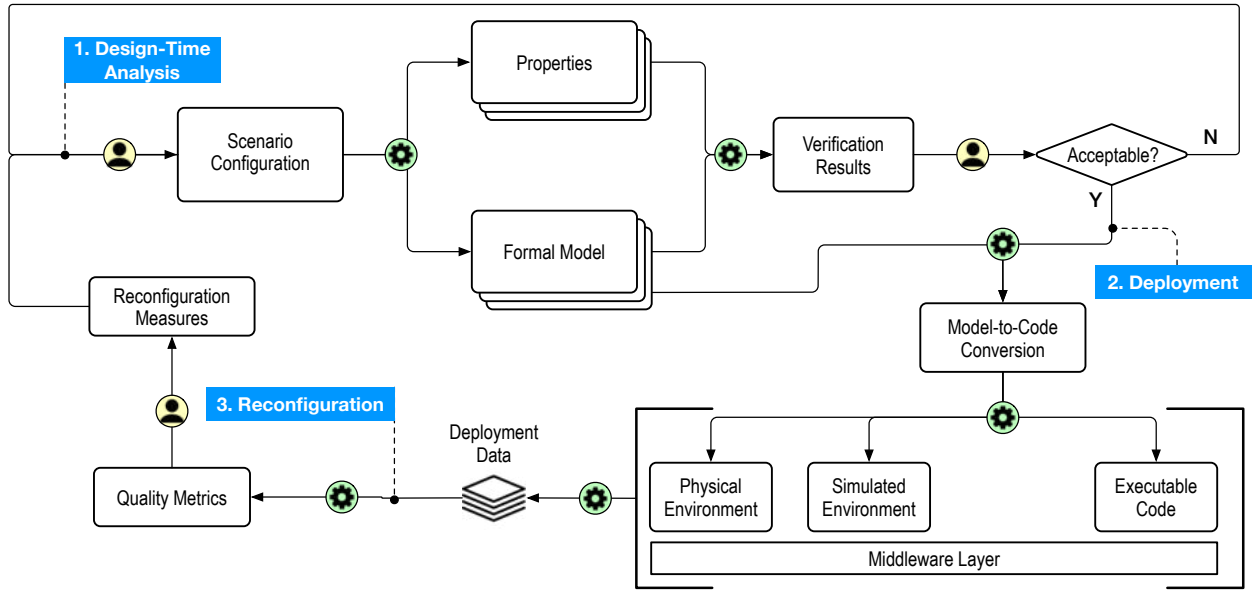


Figure 2: Diagram representing the model-driven framework operational workflow. Yellow circles mark actions performed by the *user* (i.e., the application designer) and green circles correspond to the *automated* tasks. The beginning of each phase of the framework is marked in blue and numbered according to the execution order.

- 99 2. A modeling pattern to incorporate a **stochastic** characterization of fatigue into the formal model of  
100 human behavior. In [13, 14], we presented an early version of the model formalizing an interactive robotic  
101 scenario as a Network of Hybrid Automata, with probabilistic edges capturing human choices made out of  
102 free will. This paper details the refined version of the formal model as a Network of **Stochastic Hybrid**  
103 **Automata** (SHA). Specifically, automata capturing human agents are endowed with probability  
104 distributions characterizing random parameters of the fatigue phenomenon, incorporating into the  
105 formal analysis the physiological variability between individuals belonging to different age groups or in  
106 different states of health. This source of variability (only briefly mentioned in [16]), in addition to the  
107 probabilistic characterization of human behavior, is accounted for by the formal analysis performed  
108 through **Statistical Model Checking** (SMC) via the Uppaal tool [17].
- 109 3. A model of the robot's **battery** charge and discharge dynamics refined (compared to [13, 14]) to fit  
110 the behavior of the robotic agent used for the experimental validation. The model fitting procedure  
111 increases the accuracy of the SHA modeling the battery when compared against field observations.
- 112 4. An extensive **experimental validation** to assess whether the developed formal model adheres to  
113 reality. Experimental scenarios are built by using elements that recur in 24 real-world exemplars existing  
114 in the literature and addressing service robotics in healthcare. Design-time estimations are compared  
115 with data collected by deploying benchmark applications implementing a Digital-Twin architecture.  
116 The validation activity aimed at assessing the accuracy of the formal model and the SMC results when  
117 deploying the application in a physical environment with a real robotic platform and, if necessary,  
118 with virtual human agents controlled by a real operator. We exploit a statistical technique based on  
119 Clopper-Pearson confidence intervals [18] to estimate the mission success probability range observed  
120 in reality and critically compare such results with those obtained by performing SMC experiments.  
121 Furthermore, we illustrate the application of the whole model-driven development framework to case  
122 studies capturing assistive tasks in a healthcare setting to assess how it supports the application  
123 designer from early design to reconfiguration.

124 The paper is structured as follows: Section 2 illustrates the theoretical background of the work; Section 3  
125 illustrates the design-time analysis phase of the framework, specifically introducing the conceptual model of

126 the framework’s domain and the DSL; Section 4 presents the refined SHA Network; Section 5 illustrates the  
 127 deployment and reconfiguration phases of the framework; Section 6 presents and discusses the experimental  
 128 validation results; Section 7 surveys related works in the literature; Section 8 concludes. For the interested  
 129 reader, Appendix A reports a detailed presentation of SHA semantics, while DSL specifications used for the  
 130 experimental validation are reported in Appendix B.

## 131 2. Background

132 This section illustrates the pre-existing theoretical concepts constituting the foundation of our work.  
 133 Specifically, we provide a formal definition of Stochastic Hybrid Automata (SHA) and illustrate their features  
 134 through a running example inspired by [17, Section 4]. We adopt SHA to model the robotic application, as  
 135 SHA can capture complex temporal dynamics of physical phenomena, such as the fatigue of human agents  
 136 and the battery charge/discharge for the robots, but also the digital aspects of the application, such as its  
 137 operating state and logical behavior evolving over the time.

138 **Example 2.1.** The example captures a system composed by a room with a heating system, whose model is  
 139 shown in Fig. 3a, and the thermostat controlling its temperature, shown in Fig. 3b. The room temperature  
 140 is the main physical phenomenon of the system, which is modeled by real variable  $T$  in the automaton in  
 141 Fig. 3a. The thermostat is modeled through two different operating states *on* and *off*, and it evolves by  
 142 alternating states *on*, which makes the room warmer, and *off*, thus, letting the room temperature decrease  
 143 naturally. When the thermostat is *off*, as soon as temperature  $T$  decreases below a threshold  $T_{th_1}$ , hence the  
 144 condition  $T \geq T_{th_1}$  labeling location *off* does not hold (resp., exceeds a threshold  $T_{th_2}$ , hence the condition  
 145  $T \leq T_{th_2}$  labeling location *on* does not hold), the thermostat switches the heating on (resp., off). The  
 146 triggering of the event, indicated with symbol *on!*, makes the thermostat modify its operating state, hence  
 147 moving to *on* (resp., deactivate the heating, hence moving to *off*). The room temperature is modeled in three  
 148 different situations, that are represented by the automaton in Fig. 3a: the one for which the temperature  
 149 decreases due to the absence of heating, i.e., *cool*, and two situations for which the temperature increases at  
 150 different rates, i.e., *high* and *low*, when the heating is on. The temperature grows according to differential  
 151 equation  $\dot{T} = \theta - \frac{T}{R}$  when the thermostat is on and decreases according to  $\dot{T} = -\frac{T}{R}$  when it is off, where  $R$  is  
 152 a constant and  $\theta$  is a randomly distributed parameter, i.e., whose value depends on a probability distribution.  
 153 At the onset of the system, the thermostat is off, hence the room is cooling down, and the temperature is  
 154 conventionally initialized with value  $T_{th_1}$ . This value allows the thermostat to spend a non null amount of  
 155 time in location *cool*, where constraint  $T \geq T_{th_1}$  limits the temperature inferiorly. The initialization of  $T$   
 156 is realized by the update on the edge entering location *cool*. When event *on!* (resp., *off!*) is fired by the  
 157 thermostat, the room simultaneously reacts to it; hence, both the thermostat and the room modify their  
 158 state at the same time, i.e., they synchronize. Reacting to the event, indicated with symbol *on?* (resp., *off?*),  
 159 causes the room temperature to rise and the automaton to change location to either *high* or *low* (resp., the  
 160 room temperature to decrease and the automaton to move to *cool*). The room can be heated at a high  
 161 or low rate (e.g., if a window is closed or open, respectively): the choice is made probabilistically when  
 162 the automaton synchronizes with event *on!*. The probability weights are known and amount to  $p_H$  and  $p_L$ ,  
 163 respectively. Parameter  $\theta$  in Fig. 3a is a realization of a Normal distribution with mean  $\mu_H$  and standard  
 164 deviation  $\sigma_H$  (indicated as  $\mathcal{N}(\mu_H, \sigma_H^2)$ ) when the room is heating quickly because the window is closed (the  
 165 subscript “H” stands for “high” rate of heating). Conversely, when the room is heating slowly because the  
 166 window is open, the probability distribution is  $\mathcal{N}(\mu_L, \sigma_L^2)$  (subscript “L” stands for “low” rate of heating).  
 167 Throughout the paper, we express that a random parameter  $\theta$  is a realization of random variable  $\Theta$  governed  
 168 by distribution  $\mathcal{N}(\mu, \sigma)$  through notation  $\theta \sim \mathcal{N}(\mu, \sigma)$ .

169 Thorough investigation on SHA is given in the following [19, 20, 21]. Let  $Z$  be a set of symbols; we  
 170 indicate with  $\Gamma(Z)$  the set of conjunctions of constraints of the form  $\chi_1 \sim \chi_2$ , where  $\sim$  is a relation in  $\{<, =\}$   
 171 and  $\chi_i$  ( $i \in \{1, 2\}$ ) is an arithmetical term defined by the sum of the elements in  $Z$  and  $\mathbb{N}$  (e.g.,  $z_1 + z_2 + 3$ ,  
 172 with  $z_1, z_2 \in Z$ ). By definition,  $\Gamma(Z)$  includes the logical constants **true** and **false**, defined as trivially true  
 173 formulae (e.g.,  $0 = 0$ ) or trivially false formulae (e.g.,  $0 = 1$ ). We indicate with  $\Xi(Z)$  the set of updates on

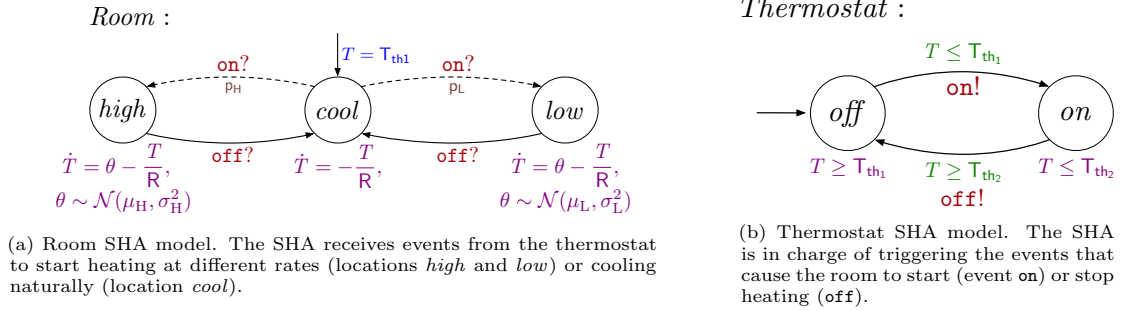


Figure 3: Example of SHA network. Dashed arrows model probabilistic transitions with weights (in brown)  $p_H$  and  $p_L$  and solid arrows represent transitions with weight 1. Flow conditions, probability distributions, and exponential rates are in purple, channels in red, and edge conditions in green, respectively.

174 elements of  $Z$ . An *update* in  $\Xi(Z)$  (for example,  $z' = z + 2$ ) is a constraint where free variables are elements  
 175 of  $Z$  (e.g.,  $z \in Z$ ) and of its primed version  $Z'$  (e.g.,  $z' \in Z'$ ). We indicate the set of non-negative real  
 176 numbers with  $\mathbb{R}_+$  and, with  $\mathbb{R}^Z$ , the set of assignments to variables of  $Z$  (i.e., *valuations*).

177 **Definition 1.** A *Stochastic Hybrid Automaton* is a tuple  $\langle L, W, \mathcal{F}, \mathcal{D}, \mathcal{I}, C, \mathcal{E}, \mu, \mathcal{P}, l_{\text{ini}} \rangle$ , where:

- 178 1.  $L$  is the set of **locations** and  $l_{\text{ini}} \in L$  is the initial location;
- 179 2.  $W$  is the set of real-valued **variables** of which clocks  $X \subseteq W$ , dense-counter variables  $V_{\text{dc}} \subseteq W$ , and  
 180 constants  $K \subseteq W$  are subsets;
- 181 3.  $\mathcal{F} : L \rightarrow \{(\mathbb{R}_+ \cup (\mathbb{R}_+ \times \mathbb{R})) \rightarrow \mathbb{R}^W\}$  is the function assigning a set of **flow conditions** to each location,  
 182 where flow conditions are (continuous) functions with one or two parameters, which assign a valuation  
 183 to every time instant in  $\mathbb{R}_+$  or to every pair constituted by a time instant in  $\mathbb{R}_+$  and a constant value  
 184 in  $\mathbb{R}$ , depending on the location;
- 185 4.  $\mathcal{D} : L \rightarrow \{\mathbb{R} \rightarrow [0, 1]\}$  is the partial function assigning a **probability distribution** from  $\{\mathbb{R} \rightarrow [0, 1]\}$   
 186 to locations which feature flow conditions with two parameters;
- 187 5.  $\mathcal{I} : L \rightarrow \Gamma(W)$  is the function assigning a (possibly empty) set of **invariants** to each location;
- 188 6.  $C$  is the set of **channels**, including the internal action  $\epsilon$ ;
- 189 7.  $\mathcal{E} \subset L \times C_{!} \times \Gamma(W) \times \wp(\Xi(W)) \times L$  is the set of **edges**, where  $C_{!} = \{c! \mid c \in C\} \cup \{c? \mid c \in C\}$  is  
 190 the set of events involving channels in  $C$ . Given an edge  $(l, c, \gamma, \xi, l') \in \mathcal{E}$ ,  $l$  (resp.  $l'$ ) is the outgoing  
 191 (resp. ingoing) location,  $c$  is the edge event,  $\gamma$  is the edge condition and  $\xi$  is the edge update. For each  
 192  $l \in L$ ,  $\mathcal{E}(l) \subseteq C_{!} \times \Gamma(W) \times \wp(\Xi(W)) \times L$  is the set of edges outgoing from  $l$  (for each  $(c, \gamma, \xi, l') \in \mathcal{E}(l)$   
 193 then  $(l, c, \gamma, \xi, l') \in \mathcal{E}$  and viceversa);
- 194 8.  $\mu : (L \times \mathbb{R}^W) \rightarrow \{\mathbb{R}_+ \rightarrow [0, 1]\}$  is the function assigning a **probability distribution** from  $\{\mathbb{R}_+ \rightarrow [0, 1]\}$   
 195 to each configuration of the SHA, where configurations are  $(l, v_{\text{var}})$  pairs constituted by a location  $l \in L$   
 196 and a valuation  $v_{\text{var}} \in \mathbb{R}^W$ ;
- 197 9.  $\mathcal{P} : L \rightarrow \{(C_{!} \times \Gamma(W) \times \wp(\Xi(W)) \times L) \rightarrow [0, 1]\}$  is the partial function assigning a discrete **probability**  
 198 **distribution** from  $\{(C_{!} \times \Gamma(W) \times \wp(\Xi(W)) \times L) \rightarrow [0, 1]\}$  to locations such that, for each  $l \in L$ ,  
 199  $\mathcal{P}(l)$  is defined if, and only if,  $\mathcal{E}(l)$  is non-empty; also, the domain of the distribution is  $\mathcal{E}(l)$  (hence,  
 200  $\sum_{\alpha=(c!, \gamma, \xi, l') \in \mathcal{E}(l)} \mathcal{P}(l)(\alpha) = 1$  holds).

201 In SHA, real-valued variables (i.e., a generalization of clocks) can evolve in time according to generic  
 202 expressions referred to as *flow conditions* [19]. The flow conditions constraining the evolution over time of

203 variables in  $W$  are defined through sets of Ordinary Differential Equations (ODEs). This feature makes  
 204 SHA a suitable formalism to model systems with complex dynamics, as it is possible to model through  
 205 flow conditions, for example, laws of physics or biochemical processes. ODEs constraining clocks (for which  
 206  $\dot{x} = 1$  holds for all  $x \in X$ ), dense-counter variables, and constants (where  $\dot{v} = 0$  holds for all  $v \in V_{dc} \cup K$ )  
 207 are special cases of flow conditions.

208 If a variable  $\theta \in V_{dc}$  is an independent term for a flow condition  $f \in \mathcal{F}(l)$  on location  $l \in L$ , i.e.,  $f = f(t, \theta)$ ,  
 209 and  $\theta$  is interpreted as a randomly distributed parameter, then  $f$  is a **stochastic process** [22]. We limit  
 210 the analysis to flow conditions depending on *at most one* random parameter, as per Definition 1, which  
 211 is enough to model human-robot interaction within the scope of our work. For example, in the SHA  
 212 shown in Fig. 3a the room temperature is modeled by real-valued variable  $T \in W$ . When temperature is  
 213 decreasing, it is constrained by the flow condition  $\dot{T}(t) = -T(t)/R$ , where function  $\dot{T}(t) \in \mathcal{F}(cool)$  depends  
 214 on time only and has solutions in  $\mathbb{R}$ . When temperature is increasing, it evolves according to flow condition  
 215  $\dot{T}(t, \theta) = \theta - T(t, \theta)/R$ , depending both on time and random parameter  $\theta$ . The domain of  $\dot{T}(t, \theta)$  is, thus,  
 216  $\mathbb{R}_+ \times \mathbb{R}$  and its solutions belong to  $\mathbb{R}$ . Interested readers are referred to Appendix A for a detailed presentation  
 217 of SHA semantics.

218 SHA are eligible for Statistical Model Checking (SMC) [23]. SMC requires a model  $M$  with stochastic  
 219 features (the SHA network), and a property  $\psi$  expressed, in our case, in Metric Interval Temporal Logic  
 220 (MITL) over atomic propositions, belonging to set AP [24], which represent, for instance, constraints over  
 221 set  $W$  (e.g.,  $w < 10$ ), or automata locations (e.g., *cool* in Fig. 3a). SMC experiments are carried out with  
 222 the Uppaal tool. Unlike exhaustive model-checking, SMC does not explore the state space but it applies  
 223 statistical techniques to a set of *traces* entailed by the formal model to estimate the *probability* of the desired  
 224 property holding. More specifically, we compute the value of expression  $\mathbb{P}_M(\psi)$  to estimate the probability of  
 225  $\psi$  holding for  $M$  [17]. Property  $\psi$ , in our framework, is of the form  $\diamond_{\leq \tau} \mathbf{ap}$ , where  $\diamond$  is the metric “eventually”  
 226 operator and  $\mathbf{ap} \in \text{AP}$ . Formula  $\diamond_{\leq \tau} \mathbf{ap}$  is true if  $\mathbf{ap}$  holds within  $\tau$  times units from time instant 0, i.e., the  
 227 onset of the system. If the value of  $\mathbb{P}_M(\psi)$  is compared against a threshold  $\vartheta \in [0, 1]$ , i.e., formula  $\mathbb{P}_M(\psi) \sim \vartheta$   
 228 is evaluated, where  $\sim \in \{\leq, \geq\}$ , the result of the SMC experiment is a Boolean value indicating whether  
 229 the probability of  $\psi$  holding for  $M$  is  $\sim$  than  $\vartheta$  (thus, **true**) or not (yielding **false**), which is calculated  
 230 through *hypothesis testing*. Otherwise, the SMC experiment returns a *confidence interval*  $[p_{\min}, p_{\max}]$  of  
 231 property  $\psi$  holding for  $M$ , with  $p_{\min}, p_{\max} \in [0, 1]$ , which is calculated according to the Clopper-Pearson  
 232 method [18]. To determine when the generated set of traces is sufficient to conclude the experiment, Uppaal  
 233 checks the length of interval  $\epsilon = (p_{\max} - p_{\min})/2$ . Uppaal stops generating new traces when  $\epsilon \leq \epsilon_{\text{th}}$  holds,  
 234 where  $\epsilon_{\text{th}}$  is an experimental parameter indicating the maximum desired estimation error. The smaller  $\epsilon_{\text{th}}$ ,  
 235 the more accurate the estimation must be and, thus, more traces are required. Similarly, by applying the  
 236 same Monte Carlo-based simulation, it is possible to calculate the expected value of distributions defined  
 237 by means of functions  $\max$  (maximum) and  $\min$  (minimum) when they are applied to stochastic processes,  
 238 such as expressions of variables in  $W$ . Given an upper bound  $\tau$ , formulae  $E_{M, \tau}[\max(v)]$  and  $E_{M, \tau}[\min(v)]$ ,  
 239 with  $v \in W$ , indicate respectively the expected value of the maximum and minimum value of variable  $v$   
 240 along executions that last *at most*  $\tau$  time units. For example, with the model of Fig. 3, we can compute the  
 241 probability of getting to operational state *high* within 10 seconds since the onset of the system by evaluating  
 242 the formula  $\mathbb{P}_M(\diamond_{\leq \tau} \mathbf{high})$  with  $\tau = 10$  and the expected value of the maximum and minimum temperature in  
 243 the room by computing  $E_{M, \tau}[\max(T)]$  and  $E_{M, \tau}[\min(T)]$ . In our framework, the SHA network  $M$  modeling  
 244 an interactive scenario is put through SMC (without specifying a probability bound  $\vartheta$ ) to estimate the  
 245 probability of success (corresponding to expression  $\mathbb{P}_M(\diamond_{\leq \tau} \mathbf{scs})$ , where Boolean variable *scs* becomes true  
 246 when the mission is completed), and to estimate the (average of the) maximum value of the fatigue of human  
 247 agents (corresponding to formula  $E_{M, \tau}[\max(F)]$ , where  $F$  is a real-value variable for the human fatigue) and  
 248 of the minimum battery charge of the robot serving in a scenario (corresponding to formula  $E_{M, \tau}[\min(C)]$ ,  
 249 where  $C$  is a real-value variable for the battery charge).

### 250 3. Design-Time Analysis of HRI Scenarios

251 This section illustrates the design-time analysis phase (phase 1 in Fig. 2), introducing the conceptual  
 252 model of the scenarios which underpins the DSL and the DSL developed to specify HRI scenarios.

As represented in Fig. 2, the design-time analysis phase begins by configuring the scenario through a custom DSL (task “Scenario Configuration” in Fig. 2). The DSL file is automatically processed to generate the SHA network and set of properties according to the user’s specifications. The designer specifies the characteristics of the robots, the involved humans, the geometrical representation of the environment layout, and the robotic missions.

Our framework features a predefined (yet extensible) set of high-level patterns identifying recurring interaction contingencies in assistive applications (e.g., a robot which follows a human). A *service* corresponds to an interaction pattern; therefore, for each service, the robot must perform the actions implied by the associated pattern (e.g., retrieve an object and deliver it back to the human). For each mission, the robot must provide services requested by the human in the order specified by the designer. Since the framework is not tied to a specific robot manufacturer or model, robotic platforms available in the fleet may not be pre-programmed to perform all required tasks. Therefore, the framework envisages an *ad-hoc* robot controller, hereinafter referred to as “*orchestrator*”, in charge of monitoring the state of the system and sending commands to the robotic agent and suggestions to human subjects in conformity with the interaction patterns. Moreover, it is paramount to take into account the robot’s level of charge and charge/discharge cycles (whose parameters vary between among different battery models), which may impact the duration of the mission (thus, its probability of success within a certain time range).

Under these premises, each mission is modeled by a SHA network featuring the following automata:

- A.  $\mathcal{A}_{h_i}$  with  $i \in [0, N_h - 1]$  modeling the  $N_h$  humans involved in the scenario;
- B.  $\mathcal{A}_r$  modeling the mobile robot;
- C.  $\mathcal{A}_b$  modeling the robot’s battery;
- D.  $\mathcal{A}_o$  modeling the orchestrator.

The tool then automatically verifies through Uppaal the specified properties. At the end of the design-time analysis phase, the designer manually examines the verification results and assesses whether they satisfy their quality criteria, for example if the probability of success is sufficiently high or the expected value of fatigue is not excessive for any human. If results are not acceptable, the designer modifies the scenario (e.g., missions, environment layout, fatigue profiles) and repeats the analysis. If results are acceptable, the application can move forward to deployment or simulation.

### 3.1. Conceptual Model of HRI Scenarios

The framework covers human-robot interaction scenarios with specific characteristics. Mainly, scenarios must take place in a known layout (thus, robotic missions carried out in unknown environments do not fall within the scope of this work) and the service sequence does not change when the application is already running. Fig. 4 shows the conceptual model (represented as a Class Diagram) for the scenarios capturing the main entities they are composed of and their relations. The diagram, described in detail in the following, constitutes the conceptual foundation of the DSL and the working assumptions that underlie the formal model. Configuring a specific scenario through the DSL to be formally verified is equivalent to defining an instance (i.e., an Object Diagram) of the conceptual model.

A **Scenario** comprises at least one robotic mission. Each **Mission** is *set in* a known **Layout**, which we represent as a *composition* of one or multiple two-dimensional rectangular areas. Each **Area** is a *composition* of four corner **Points**, each characterized by a pair of Cartesian coordinates  $x$  and  $y$ . A **Layout** also includes a relevant subset of points, called **Points Of Interest (POI)**, that can be the target of an action, such as room entrances, cupboards, and the robot’s recharge station. Missions, areas and POIs are identified through attribute **name**.

A mission is a sequence of services *requested* by a human and *provided* by a robot (specifically, humans are served according to their **id**). Each **Service** conforms to an interaction pattern (attribute **ptrn**) and has a target **POI**. Patterns (i.e., the items of enumeration **InteractionPattern**) group common interaction contingencies and are listed in the following:





317 kit during an emergency). Both agents move to the location of the resource (captured by attribute  
 318 **target**) to reach it as quickly as possible. The competition ends when either of the agents reaches the  
 319 target location (effectively *winning* the competition). The human may autonomously decide to stop  
 320 walking at any time.

321 P5. **HumanRescuer**: the pattern captures the robot requiring human intervention to complete a task, such  
 322 as pressing a button to call the elevator or opening a closed door. In this case, the robot will emit  
 323 audible or visible signals to notify its need for human support. The human autonomously decides to  
 324 support the robot, move to the robot’s current location (captured by attribute **target**), perform the  
 325 required action and conclude the interaction.

326 P6. **HumanApplicant**: the pattern captures the human requiring the robot’s support in performing a  
 327 certain task that implies timely or close-contact interaction, such as feeding a patient or administering  
 328 medication. In this case, as soon as the service starts, the human waits for the robot to approach  
 329 their current location (attribute **target**). When the robot is sufficiently close, the action requiring  
 330 synchronization starts. The human may autonomously decide to interrupt the action and resume at  
 331 any time.

332 Agents enact the mission. Abstract class *Agent* has a **name**, **id**, and starting position **start** within the  
 333 layout. In case of human agents, the **id** attribute determines the order in which humans are served. In case  
 334 of robotic agents, the **id** attribute determines the order in which missions are assigned to robots in the fleet  
 335 in case of multi-robot missions [16]. Agents are endowed with sensors that share a new reading every  $T_{\text{poll}}$   
 336 instants. Within the scope of our framework, there are two possible *specializations* of an *Agent*: humans  
 337 and robots. For each robot, attribute **type** from enumeration **RobotType** defines its commercial model (e.g.,  
 338 “TurtleBot3” or “Tiago”). We assume that a **Robot** moves with a trapezoidal velocity profile, whose maximum  
 339 acceleration  $a_{\text{max}}$ , linear velocity  $v_{\text{max}}$ , and angular velocity  $\omega_r$  are *derived* from attribute **type**. Each **Robot** is  
 340 *powered by* a lithium **Battery** with initial charge  $C_0$ . Class **Battery**’s attribute  $C_{\text{fail}}$  corresponds to the lowest  
 341 voltage under which the device must not move to prevent the battery pack from being damaged.

342 SHA modeling human behavior include a model of physical fatigue. Each **Human** has a  $p_f$  attribute  
 343 determining their fatigue profile (see the **FatigueProfile** enumeration in Fig. 4), which determines their  
 344 proneness to fatigue and recovery based on physiological factors. We distinguish subjects by age (**Young/Elderly**)  
 345 and state of health (**Healthy/Sick**) or whether they are affected by a severe respiratory syndrome that hinders  
 346 their ability of deambulation (**SARSPatient**), obtaining five possible fatigue profiles. Attribute **v** specifies  
 347 the average walking speed. Since human behavior is unpredictable in a real setting, our model includes a  
 348 probabilistic approximation of human haphazard behavior (e.g., the possibility to ignore a robot’s instruction  
 349 or start and stop freely during the interaction). Therefore, a **Human** also features attribute  $p_{\text{fw}}$  from which  
 350 attributes **obey**,  $\text{FW}_{\text{max}}$ , and  $\text{FW}_{\text{th}}$  determining the probability with which such behavior manifests itself are  
 351 derived. The  $p_{\text{fw}}$  attribute has four possible values (corresponding to the elements of the **FreeWillProfile**  
 352 enumeration): **Normal**, **High**, **Low**, or **Disabled**. The latter results in human free will being entirely ignored at  
 353 design-time, which may only be used for a preliminary test of the scenario setup.

354 Agents and batteries are equipped with sensors that during deployment share data with the orchestrator  
 355 over dedicated topics handled by the middleware layer (based on ROS, see Fig. 2)[41]. The SHA network  
 356 features a model of ROS publisher nodes (i.e., instances of class **ROSPubNode** in Fig. 4) mimicking the delay  
 357 with which messages are processed and published. These delays are normally distributed with mean  $l_{\text{mean}}$   
 358 and variance  $l_{\text{var}}$  [25].

359 The **Orchestrator** monitors the state of the system by *subscribing* to sensor readings’ topics of the human’s  
 360 position, fatigue, robot’s position and battery charge. The **Orchestrator** periodically checks the state of the  
 361 system against its policies every  $T_{\text{int}}$  time units and processes data for  $T_{\text{proc}}$  time units. While processing, it  
 362 checks sensor-collected data against a set of thresholds:  $D_{\text{stop}}$  and  $D_{\text{restart}}$  determine the human-robot distance  
 363 that causes the robot to stop and wait or restart, respectively;  $C_{\text{rech}}$  and  $C_{\text{restart}}$  correspond to the battery  
 364 charge levels that cause the robot to start or stop recharging, respectively;  $F_{\text{stop}}$  and  $F_{\text{restart}}$  correspond to the  
 365 human fatigue levels inducing the human to stop and rest or resume the action, respectively.

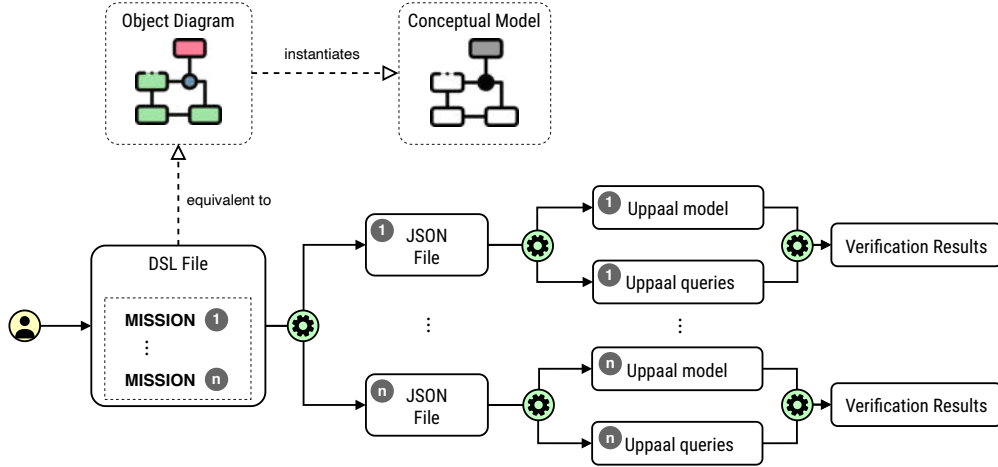


Figure 5: Diagram representing the process that translates a DSL file into Uppaal models. Solid arrows represent operational tasks while dashed arrows represent conceptual equivalences. Solid arrows are marked to distinguish the actions performed by the user from those performed automatically. Boxes are numbered to identify the relations between defined missions and the output files of each phase.

Finally, for each scenario, the analyst is interested in quality metrics to be computed referred to as “queries”. The framework currently supports three query types (i.e., probability calculation, expected value calculation, and generation of traces) supported by Uppaal. However, different queries and tools can easily be embraced. For each Query, it is necessary to specify its type, time bound  $\tau$  and—optionally—the maximum number of system traces generated to verify the property, i.e., attribute  $R$ . We remark that  $R$  should only be specified to limit performance issues during preliminary testing while it is normally advisable to let the verification tool compute the number of runs necessary to perform verification with the required confidence level. The different query types (modeled by enumeration `QueryType`) allow the designer to estimate:

- Q1. the probability of the mission ending with success (item `P_SCS`) within the time bound. Success occurs when all services have been completed;
- Q2. the probability of the mission ending in failure (item `P_FAIL`) within the time bound. Failure occurs either when the robot is fully discharged and cannot move autonomously or the human is fully fatigued (note that the mission not ending in success due to an insufficient time bound does not constitute a failure, thus the results of a `P_FAIL` and `P_SCS` query do not necessarily sum to 1);
- Q3. the expected maximum value of fatigue (item `E_FTG`) for all humans within the time bound;
- Q4. the expected minimum battery charge (item `E_CHG`) value within the time bound;
- Q5. one or multiple (i.e., specified as parameter  $R$ ) system traces to have a more detailed overview of how the system behaves during the execution of the mission (item `SIM`).

### 3.2. Domain-Specific Language

As represented in Fig. 5, configuring a scenario through the DSL is semantically equivalent to defining an Object Diagram of the conceptual model in Fig. 4. Therefore, the developed DSL features primitives allowing for the creation of instances of concrete scenarios that reflect the conceptual model. Each primitive is presented in the upcoming subsections through an example.

Fig. 5 shows how DSL files are converted into SMC experiment instances. Each DSL file defines a single scenario, which includes the layout geometry, the points of interest, the agents, and the mission, and represents a well-formed DSL model if specific properties are met (e.g., rooms have non-null area, agents are located within the boundaries of the environment, etc.). Well-formedness properties are automatically

393 verified by the translator every time a DSL file undergoes the conversion process. Each mission in the DSL  
 394 model is subject to formal verification separately, requiring, thus, a separate formal model. The conversion  
 395 process features an intermediate phase: a JSON file containing the mission’s characteristics is generated  
 396 for each mission. Each JSON file is then converted into two files, one with the Uppaal model and one with  
 397 the queries to perform the SMC experiment. The intermediate JSON notation, which is a lightweight and  
 398 well-established standard, decouples the DSL from the specific verification tool and makes the framework  
 399 flexible to the introduction of different verification tools or different DSLs. JSON files are also exploited to  
 400 automatically set up the deployment environment.

401 We illustrate the DSL features and how they can be exploited to model the illustrative scenario in Section  
 402 1.1. The DSL does not have a specific statement for objects of class `Scenario` because each file inherently  
 403 instantiates a single scenario, possibly including several missions. Every mission in a scenario consists of  
 404 four independent sections, each one identified by the keyword `define`, concerning: layout definition, list of  
 405 agents in the scene, list of services, and list of queries to be computed. Orchestrator and ROS nodes are  
 406 instantiated automatically when the specification is translated into the model to be used for verification.

### 407 3.2.1. Layout, Areas and POIs

While modeling the HRI scenario, the user must specify the layout where the agents will operate. The DSL allows users to model different layouts (such as different building floors or different sections of the same floor) through statement:

`define layout`

408 which includes a non-empty list of areas and POIs.

The DSL captures all layouts made up of adjacent rectangular areas (i.e., it does not capture curved or diagonal walls), each defined as:

`area id in (x1, y1) (x2, y2)`

409 where coordinate pairs  $(x_1, y_1)$  and  $(x_2, y_2)$  define one of the area’s diagonal segments from which the other  
 410 two corner points are automatically inferred upon generating the SHA network to save manual effort on the  
 411 user’s side. Areas’ corners are validated to ensure that they correctly identify a diagonal, i.e., that  $x_1 \neq x_2$   
 412 and  $y_1 \neq y_2$  hold. Layout-related declarations are validated to check whether there are disconnected areas  
 413 (i.e., all areas must be *reachable* from any point in the layout) and that no pair of areas overlap entirely.

POIs with their coordinates are declared through the following statement:

`poi id in (x, y)`

414 Although verification tools only handle adimensional variables, the DSL requires the specification of the length measurement units to ensure that the layout size is consistent with the robot’s speed. The measurement unit is specified through the following statement, where  $m_u \in \{\text{km}, \text{m}, \text{cm}\}$ :

`param measurement_unit mu`

415 The specification of the layout in Fig. 1 is given in Listing 1: the layout features two areas (`a1` and `a2`)  
 416 and three POIs corresponding to the recharge station (`RC`), the waiting room entrance (`WR`), `KIT1`, and  
 418 `KIT2`. All coordinates are expressed in meters (`m`), consistently with Fig. 1b.

### 419 3.2.2. Agents

Each mission must feature a mobile robot and at least one human requiring assistance, specified in two independent sections that are identified through keywords `define robots` and `define humans`, respectively. The DSL allows designers to declare each available robot through statement:

`robot name in (x, y) id id type type charge C0`

---

**Listing 1** DSL section defining layout areas and POIs.

```
1 param measurement_unit m
2 define layout:
3   area A1 in (0.0, 17.5) (40.0, 7.5)
4   area A2 in (40.0, 25.0) (50.0, 0.0)
5   poi RC in (25.0, 17.5)
6   poi WR in (49.5, 12.5)
7   poi KIT1 in (40.5, 21.25)
8   poi KIT2 in (40.5, 3.75)
```

---

**Listing 2** DSL section defining the agents and their features.

```
1 define robots:
2   robot ROB1 in (10.0, 12.5) id 1 type tiago charge 40
3   robot ROB2 in (45.0, 3.5) id 2 type turtlebot3_wafflepi charge 90
4
5 define humans:
6   human HUM1 in (5.0, 12.5) id 1 speed 80 is young_sick freewill low
7   human HUM2 in (35.0, 9.0) id 2 speed 100 is elderly_healthy freewill normal
```

---

420 where parameters `name` and `id` univocally identify the robot,  $C_0$  defines its initial level of charge, and  
421 coordinates  $(x, y)$  define its starting position. As per Section 3.1, while generating the formal model, the  
422 robot’s type determines the translational and rotational speeds, and the acceleration (attributes  $v_{\max}$ ,  $\omega_r$ , and  
423  $a_{\max}$ ) in Fig. 4) based on the model’s technical specifications. This feature of the DSL saves non-technical  
424 users the effort of retrieving these data when they might be more familiar with the type of robot available in  
425 the facility.

426 Each human is declared through the following statement:

$$\text{human name in } (x, y) \text{ id id speed } v \text{ is } p_f \text{ freewill } p_{fw} \quad (1)$$

427 where the `name` univocally identifies the human and the `id` determines the serving order (thus, it is also  
428 required to be unique). Coordinates  $(x, y)$  determine each human’s starting location. Parameters  $v$ ,  $p_f$ , and  
429  $p_{fw}$  define the walking speed, fatigue and free will profiles as described in Section 3.1. The user chooses the  
430 values of  $p_f$  and  $p_{fw}$  out of a pre-determined list, corresponding to the enumerations in Fig. 4.

431 Both the robot and human declaration blocks are validated to ensure that no pair of agents share the same  
432 `id` (within the same *Agent* generalization) nor the same `name` (also across different *Agent* generalizations). The  
433 DSL is developed under the simplifying hypothesis that agents occupy a single point in space (corresponding  
434 to their center of gravity). This modeling choice is dictated by the need to keep the DSL (and, thus, the  
435 formal model) as simple as feasible and spare the designer from defining the three-dimensional envelope of  
436 the agents’ bodies. The framework assumes that refined collision avoidance routines are already implemented  
437 at a lower level within the robotic platform and the DSL validator only checks that no pair of agents have  
438 the same center of gravity (i.e., coordinates  $(x, y)$ ).

439 The agents from the running example are defined as per Listing 2. There are two robots available (ROB1  
440 and ROB2) of different types (thus, they will have different speeds), and two humans (HUM1 and HUM2), of  
441 which one has a Young/Sick fatigue profile and low free will, whereas the second one is Elderly/Healthy and  
442 has normal free will profile.

### 443 3.2.3. Missions

Designers can declare multiple **missions** and associate them with a layout and a set of agents. Each mission is assigned to a single robot and verification experiments resulting from each mission declaration (see

---

**Listing 3** DSL section defining the mission (i.e., the sequence of services).

```
1 define mission m1 for ROB1:  
2   do robot_leader for HUM1 with target WR  
3   do robot_transporter for HUM2 with target KIT2  
4  
5 define mission m2 for ROB2:  
6   do robot_follower for HUM2 with target KIT1  
7   do robot_leader for HUM1 with target WR
```

---

**Listing 4** DSL section defining the set of queries.

```
1 define queries of mission m1:  
2   compute probability_of_success with duration 120 runs 300  
3   compute expected_fatigue with duration 120 runs 50  
4  
5 define queries of mission m2:  
6   compute probability_of_success with duration 100 runs auto  
7   compute probability_of_failure with duration 100 runs auto  
8
```

---

Fig. 5) are performed separately. A mission is declared as in the following, where parameter  $m$  is the name of the mission and  $r$  is the name of the robot it is assigned to (association *provides* in Fig. 4).

**define mission  $m$  for  $r$**

444

As described in Section 3.1, each mission consists of a sequence of services and each service adheres to one of the interaction patterns described in Section 3.1. As per Statement (1) and the conceptual model presented in Section 3.1, humans are declared independently of the interaction pattern, which is specified when declaring the service as in the following:

**do ptrn for  $h$  with target poi**

445

where *ptrn* can be either *robot\_leader*, *robot\_follower*, *robot\_transporter*, *robot\_competitor*, *robot\_applicant*, *robot\_rescuer* (corresponding to the *HumanFollower*, *HumanLeader*, *HumanRecipient*, *HumanCompetitor*, *HumanRescuer*, and *HumanApplicant* patterns, respectively),  $h$  is the name of the human requesting the service (association *requests* in Fig. 4), and *poi* instantiates attribute *target* in Fig. 4. Each service declaration is validated to ensure that  $h$  and *poi* refer to existing human agents and POIs.

450

The two missions associated with the running example are declared as in Listing 3. In mission *m1*, *ROB2* has to lead *HUM1* to POI *WR* and then deliver *KIT2* to *HUM2*. In *m2*, *ROB2* has to follow *HUM2* to *KIT1*, then lead *HUM1* to *WR*.

452

#### 453 3.2.4. Queries

Finally, the designer specifies which experiments to perform for each mission. The DSL captures the set of queries in Fig. 4 and described in Section 3.1. A query is declared through the following statement:

**compute query with duration  $\tau$  runs  $R$**

454

where *query* can be either *probability\_of\_success*, *probability\_of\_failure*, *expected\_charge*, *expected\_fatigue*, or *simulation*. Parameter  $\tau$  corresponds to the time bound, while  $R$  is the bound on the number of traces

455

456 generated for the SMC experiment, whose value must be set to `auto` if the user wants the verification tool to  
 457 compute the required number of runs.

458 A possible set of queries for missions 1 and 2 from the running example is given in Listing 4: the SMC  
 459 experiments will estimate the probability of success and maximum fatigue value for all humans for `m1`, and  
 460 probabilities of failure and success (with no bound on runs) for `m2`.

#### 461 4. Formal Modeling HRI with Uncertain Human Behaviors

462 In this section, we illustrate the modeling approach we have adopted to *map* aspects of the real system to  
 463 SHA features. Subsequently, we present in more detail the automata constituting the SHA network, i.e., the  
 464 humans, the robot and its battery, and the orchestrator.

465 The high-level goal of the SHA network is to capture the agents’ behavior based on their current *operating*  
 466 *state* (e.g., the human resting or walking). For every agent in the scenario and automaton  $\mathcal{A}$  modeling  
 467 its behavior, defined as in Definition 1, every operating state of the agent corresponds to a *location* in  $L$ .  
 468 SHA capture the evolution of relevant quantitative attributes of the real system, such as human fatigue  
 469 and battery level of charge. Each physical attribute, characterizing a human or a component, corresponds  
 470 to a real-valued variable in set  $W \setminus \{X \cup V_{dc} \cup K\}$  of their modeling automata and flow conditions  $\mathcal{F}(l)$ ,  
 471 associated with a location  $l$ , reproduce the set of ODEs constraining the evolution of real-valued variables in  
 472 that specific operating state.

473 The switch between two operating states consists of an *edge* between the two corresponding locations.  
 474 Recurrent features of the specific systems that our modeling approach targets identify two *types* of switches,  
 475 which we refer to as *controllable* or *uncontrollable*. We remark that we use these two terms in a manner that  
 476 is specific to our framework, and they are not part of the standard terminology of the formalism (for example,  
 477 they are unrelated to the notion of controllable and uncontrollable edges in Timed Game Automata [26]);  
 478 instead, they are merely aliases for specific edges recurring in our SHA (i.e., subsets of  $\mathcal{E}$  from Definition 1).  
 479 A controllable switch occurs if, and only if, a specific event fires and a synchronization among two or more  
 480 automata occurs: for example, the robot starts accelerating when the orchestrator issues the command to  
 481 start moving. For this reason, all the edges modeling a controllable switch are defined with a channel in  
 482  $C \setminus \{\epsilon\}$ . Conversely, uncontrollable switches occur “naturally” in the original system due to the evolution of  
 483 the physical variables at play: for example, the human unavoidably stops moving when their fatigue level  
 484 reaches the maximum endurable threshold. Therefore, they are defined with event  $\epsilon!$ , i.e., by means of the  
 485 internal action. In every automaton of the networks capturing the targeted systems, a location  $l$  with an  
 486 outgoing edge modeling an uncontrollable switch is endowed with a set of invariants of the form  $w \leq k_2$ ,  
 487 where  $w \in W$  is the real-valued variable subject to the constraint and  $k_2 \in K$  is its maximum allowable  
 488 value (e.g.,  $F \leq 1$ , with  $w = F$  and  $k_2 = 1$ , constrains the value of the human fatigue to be less or equal than  
 489 1). The outgoing edge has condition  $w \geq k_1$ , such that  $k_1 \in K$  and  $k_2 \geq k_1$  hold. If  $k_2 > k_1$  holds, the edge  
 490 fires with probability distributed uniformly over interval  $[k_1, k_2]$ , as explained in Appendix A. If  $k_1 = k_2$   
 491 holds, the edge fires with probability 1 when  $w = k_1 = k_2$  holds (e.g., the edge from *on* to *off* in Fig. 3 where  
 492  $w = T$  and  $k_1 = k_2 = T_{th_2}$ ).

493 Since the orchestrator controls the robots by using the digital observations made by sensors on humans and  
 494 robots, the SHA modeling physical dynamics (i.e., humans and robots) feature dense-counter variables (set  
 495  $V_{dc}$ ) as the discrete (i.e., digital) equivalents of real-valued variables. Dense counters are periodically updated  
 496 every  $T_{poll} \in K$  time units (where  $T_{poll}$  corresponds to the refresh period of the specific sensor) through  
 497 updates in  $\Xi(W)$  that are compatible with the ODEs modeling the dynamics of the physical attributes in  
 498 each location. To this end, every SHA in the network uses a clock  $t_{upd} \in X$  to measure the time elapsed  
 499 between two consecutive measurements and to trigger an update. Therefore, when  $t_{upd} = T_{poll}$  holds for  
 500 an automaton  $\mathcal{A}$ , hence when time  $T_{poll}$  has elapsed since the last measurement, then  $\mathcal{A}$  uncontrollably  
 501 switches to a *committed* location. A committed location is equivalent to an ordinary location with invariant  
 502  $t \leq 0$  and all incoming edges with update  $t = 0$  for some  $t \in X$ : therefore, time cannot elapse while in these  
 503 locations [27]. In that location, the dense counters modeling the latest sensor readings are immediately  
 504 notified to the orchestrator by firing an event over a dedicated channel that triggers the publishing routine  
 505 (the corresponding modeling pattern is described in detail in [15, Section IV]).

In the forthcoming sections, we present the automata constituting the SHA network, i.e., the humans, the robot and its battery, and the orchestrator, primarily focusing on the human behaviors modeled by the patterns Human Follower, Human Leader, and Human Recipient, which this work extends with respect to [13, 14] with a stochastic characterization of physical fatigue, and a factorization of the common modeling pattern capturing the periodic sensor reading update. Comparable SHA for the Human Applicant, Human Rescuer, and Human Competitor patterns are introduced in [16]. We also present a refined model of the robot’s battery, which fits the real platform used for the experimental validation. The robot-modeling SHA is presented in detail in [15] and briefly described here to preserve the self-containedness of this paper. Moreover, we report the high-level structure of the orchestrator (introduced in [13] and extended in [16]) with refined mission-management policies.

Controllable switches realize the interactions among the automata and are obtained by means of synchronization channels. Since the orchestrator implements the control logic that governs the agents, the issuing of a command by the orchestrator is modeled through a synchronization between the orchestrator and the automaton modeling the agent that reacts to the command. Hence, all the channels in (the automaton modeling) the orchestrator are labeled with `!`, whereas (the automata modeling) the humans and the battery are defined with edges having the channels labeled with `?`. Hence, The modeling has been realized by considering one single robot serving one individual at a time, even if several agents (i.e., humans and robots) can participate in the scenario. In fact, each pattern models the interaction between one pair of agents, a robot and a human, and missions are finite sequences of interactions. Since there is always only one active robot and a single served human at a time, channels representing synchronizing events between the orchestrator and the agents are not specific to a single instance of a human or a robot. A dense counter in the orchestrator, with finite domain, identifies the human currently interacting with the robot and it is evaluated by every automata modeling the humans to allow or deny the firing of the synchronization events with the orchestrator. In particular, all the edges in the automaton modeling a human agent include a condition which evaluates to true when the dense counter indicating the currently served human is equal to the value `id` uniquely identifying it (see Section 3.1) and, hence, the automaton. Moreover, even if the modeling of multiple robots serving multiple humans simultaneously is possible in theory, adding this feature would cause the models to increase in complexity. The information flow that the orchestrator realizes by issuing commands to agents, through events via channels, is as follows. The orchestrator

- informs the human to start or stop walking via channels `cmd_h_start` and `cmd_h_stop`;
- makes the robot move or stop via channels `cmd_r_start` and `cmd_r_stop`. Moreover, it starts the battery charging through channel `cmd_b_start` and interrupts the charging with channel `cmd_b_stop`, hence restoring the robot back to the mission-defined interaction.

#### 4.1. Human-Robot Interaction Patterns Model

In all interaction patterns, the SHA modeling humans differentiate between operating states based on how fatigue evolves (i.e., whether individuals are recovering or not) and how humans are interacting with the robot (e.g., they are leading the action or waiting for a robot’s action). Hence, all SHA modeling humans feature real-valued variable  $F \in W$ , capturing physical fatigue, and a dense counter  $f \in V_{dc}$  capturing the digital counterpart of  $F$ . Besides physical fatigue, for each human, suitable sensors also periodically refresh their position within the building. The position is modeled by dense counters  $h_{pos_x}$  and  $h_{pos_y}$  capturing a pair of Cartesian coordinates. Therefore, the portion of SHA modeling the update of periodic sensor readings is present in all the operating states of the human, generically indicated as `op`, and is hereinafter referred to as `⟨op⟩_pubh`. In the following sections, for a clock  $t \in X$ , we use notation  $\{t\}$  to represent update  $t = 0$  (e.g., we write  $\{t_{upd}\}$  instead of  $t_{upd} = 0$ ).

The `⟨op⟩_pubh` pattern is shown in Fig. 6. In the following, we use label `op` when describing the high-level structure of the pattern, while it is replaced by descriptive labels when referring to a specific instance of the pattern (e.g., `⟨stand⟩_pubh` and `⟨walk⟩_pubh`). The automaton features three locations: an ordinary location  $h_{\langle op \rangle}$  and two committed locations  $h_{pub_1}$  and  $h_{pub_2}$ . Location  $h_{\langle op \rangle}$  captures the human’s behavior while in a specific state (e.g.,  $h_{\langle stand \rangle}$  and  $h_{\langle walk \rangle}$ ). To this end,  $h_{\langle op \rangle}$  is endowed with invariants  $\mathcal{I}(h_{\langle op \rangle})$ , flow conditions



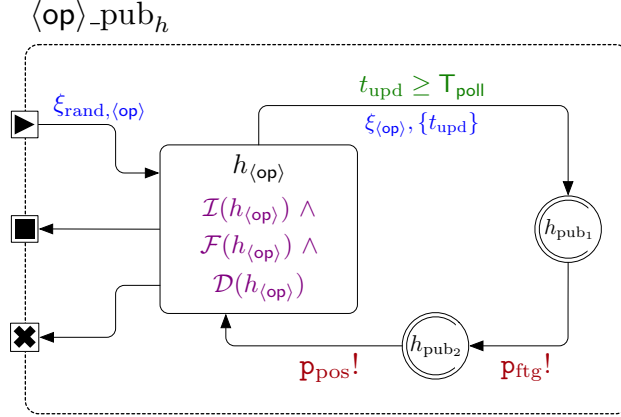


Figure 6: SHA modeling the  $\langle \text{op} \rangle\text{-pub}_h$  pattern, color-coded as in Fig. 3. Ports are marked by symbols “▶”, “■”, and “×”.

555  $\mathcal{F}(h_{(\text{op})})$ , and probability distributions  $\mathcal{D}(h_{(\text{op})})$ . For all instances of  $\langle \text{op} \rangle\text{-pub}_h$ ,  $(t_{\text{upd}} \leq T_{\text{poll}}) \in \mathcal{I}(h_{(\text{op})})$   
 556 holds. The combination of this invariant with condition  $t_{\text{upd}} \geq T_{\text{poll}}$  on the edge to  $h_{\text{pub1}}$  forces the SHA to  
 557 switch to the committed location when  $t_{\text{upd}} = T_{\text{poll}}$  holds. Upon switching, the set of updates  $\xi_{(\text{op})} \subset \Xi(W)$   
 558 (e.g.,  $\xi_{(\text{stand})}$  and  $\xi_{(\text{walk})}$ ) updates dense counters  $f$ ,  $h_{\text{pos}_x}$ , and  $h_{\text{pos}_y}$ . The effect of  $\xi_{(\text{op})}$  varies depending on  
 559 the specific state of the human. Since  $h_{\text{pub1}}$  and  $h_{\text{pub2}}$  are committed, the new values are immediately shared  
 560 with the orchestrator by firing an event through channels  $\mathbf{p}_{\text{ftg}}$  first and  $\mathbf{p}_{\text{pos}}$  right after.

561 Edges entering and leaving the  $\langle \text{op} \rangle\text{-pub}_h$  pattern are represented through *ports* (coherently with [15, 16]).  
 562 Ports are not part of the formalism, but a visualization expedient for the edges entering and leaving the  
 563 sub-automaton (arrows in and out of a port constitute the *same* transition). SHA *enter* a submachine  
 564 through the port marked by symbol “▶” (i.e., *start*) and *leave* a submachine through ports marked by  
 565 symbols “■” (i.e., *end*) and “×” (i.e., *fail*), indicating whether the operating state  $\text{op}$  ended (or stopped  
 566 momentarily) or the entire mission ended with failure (e.g., because the human is too fatigued), respectively.

567 Edge conditions, channels, and updates characterizing edges through ports vary depending on the specific  
 568  $\langle \text{op} \rangle\text{-pub}_h$  instance. The only exception is update  $\xi_{\text{rand},(\text{op})}$  on the edge through the *start* port. Update  
 569  $\xi_{\text{rand},(\text{op})}$  is featured by *all* instances of  $\langle \text{op} \rangle\text{-pub}_h$  since it determines the stochastic properties of human  
 570 fatigue when a human behaves while in a specific operating state  $\langle \text{op} \rangle$  and the way these properties are  
 571 determined is the same for every instance of  $\langle \text{op} \rangle\text{-pub}_h$ .

572 Human fatigue is a complex phenomenon driven by a wide range of factors: our approach focuses on  
 573 muscular fatigue due to physical strain. As discussed by Liu et al. [28], a muscle can be seen as a reservoir  
 574 of motor units. When physical exertion is required, motor units progressively activate and eventually cause  
 575 *fatigue* due to biochemical processes. The muscle can, then, *recover* from fatigue if it is put to rest [28, 29].  
 576 Our approach exploits the model proposed by Konz [29, 30], described by Eq.2, for which human action  
 577 undergoes alternate fatigue and recovery cycles, each one modeled by an exponential function. Fatigue and  
 578 recovery are expressed by means of function parameters, called fatigue rates, which depend on several factors  
 579 such as the age of the subject that the model represents, their health condition, etc. Each cycle is associated  
 580 with an index  $i$  uniquely identified, given time  $t$ , by function  $j : \mathbb{R}_+ \rightarrow \mathbb{N}$  (thus,  $i = j(t)$  holds). We indicate  
 581 the timestamp at which cycle  $i$  ends by  $\mathbf{t}_i$ . During both fatigue and recovery, fatigue  $F(t)$  for cycle  $i$  depends  
 582 on the residual value  $F(\mathbf{t}_{i-1})$  from the previous cycle ended at time  $\mathbf{t}_{i-1}$ . Parameters  $\lambda_i$  and  $\rho_i$  are the  
 583 fatigue and recovery rates for cycle  $i$ .

$$F(t) = \begin{cases} 1 - (1 - F(\mathbf{t}_{i-1})) \cdot e^{-\lambda_i(t-\mathbf{t}_{i-1})} & \text{(fatigue)} \\ F(\mathbf{t}_{i-1}) \cdot e^{-\rho_i(t-\mathbf{t}_{i-1})} & \text{(recovery)} \end{cases} \quad (2)$$

584 Full recovery occurs when  $F(t) = 0$  holds, whereas condition  $F(t) = 1$  models the case in which the muscle  
 585 has reached the maximum level of *endurance*. Liu et al. [28] argue that the fatigue  $F(t)$  can be seen  
 586 as ratio  $M_F(t)/M_0$ , where  $M_0$  is the total amount of motor units, and  $M_F(t)$  is the amount of fatigued

587 units at time  $t$ . Therefore,  $F(t) = M_F(t)/M_0 = 1$  holds when *every* unit composing a muscle is fatigued.  
 588 Running experiments on a pool of subjects have shown how a Normal distribution is a good fit to capture  
 589 the variability of rates  $\lambda_i$  and  $\rho_i$  in the fatigue model [31]. Furthermore, the variability of the fatigue rates  
 590 for an individual subject between different exertion cycles has been observed in [32]. The SHA modeling  
 591 the human in a scenario embeds this variability by means of probability distributions, as the automaton is  
 592 not representative for a single specific individual, but it represents a set of subjects with similar physical  
 593 characteristics. Therefore, we approximate the complexity of the fatigue phenomenon by considering each  $\lambda_i$   
 594 (resp.,  $\rho_i$ ) as a sample of distribution  $\mathcal{N}(\mu_\lambda, \sigma_\lambda^2)$  (resp.,  $\mathcal{N}(\mu_\rho, \sigma_\rho^2)$ ), whose mean and variance depend on the  
 595 fatigue profile that characterizes the class of humans under analysis.

596 By construction, every operating state of a human agent is associated with a specific fatigue profile, i.e.,  
 597 it is either a fatigue state or a recovery state. Hence, for every instance of  $\langle \text{op} \rangle$ -pub<sub>h</sub> function  $\mathcal{D}(h_{\langle \text{op} \rangle})$   
 598 is defined. Upon entering an  $\langle \text{op} \rangle$ -pub<sub>h</sub>, update  $\xi_{\text{rand}, \langle \text{op} \rangle}$  computes the fatigue/recovery rate to be considered  
 599 while the automaton is in location  $h_{\langle \text{op} \rangle}$ . To this end, every automaton modeling a human features two dense  
 600 counters  $\lambda, \rho \in V_{\text{dc}}$ , which store the current fatigue/recovery rates. Every time update  $\xi_{\text{rand}, \langle \text{op} \rangle}$  is executed,  
 601 it generates a new sample of  $\mathcal{D}(h_{\langle \text{op} \rangle})$  and assigns it to  $\rho$ , if  $h_{\langle \text{op} \rangle}$  is a recovery state, otherwise to  $\lambda$ . The  
 602 sample is generated through the Box-Müller algorithm [17]. Update  $\xi_{\text{rand}, \langle \text{op} \rangle}$  is given in Eq.3, where rate  
 603 equals  $\lambda$  if  $h_{\langle \text{op} \rangle}$  is a fatigue state and  $\rho$  otherwise;  $\mu_{\langle \text{op} \rangle}$  and  $\sigma_{\langle \text{op} \rangle}$  are the mean and standard deviation of  
 604  $\mathcal{D}(h_{\langle \text{op} \rangle})$ ;  $u_1$  and  $u_2$  are independent realizations of uniform distribution  $\mathcal{U}(0, 1)$ .

$$\xi_{\text{rand}, \langle \text{op} \rangle} : \text{rate} = \mu_{\langle \text{op} \rangle} + \sigma_{\langle \text{op} \rangle} \sqrt{-2 \ln(u_2)} \cos(2\pi u_1) \quad (3)$$

605 The values of  $\rho$  and  $\lambda$  determine the temporal evolution of the real-valued variable  $F$  and its digital counterpart  
 606  $f$  while the automaton is in  $h_{\langle \text{op} \rangle}$ . In a given operating state  $\langle \text{op} \rangle$ , if fatigue increases, flow condition  $\mathcal{F}(h_{\langle \text{op} \rangle})$   
 607 corresponds to the derivative of Eq.2(fatigue), indicated as  $f_{\text{ftg}}$  in Eq.4; otherwise, fatigue decreases and  
 608  $\mathcal{F}(h_{\langle \text{op} \rangle})$  is equal to the derivative of Eq.2(recovery), indicated as  $f_{\text{rec}}$  in Eq.5. Both equations depend on two  
 609 terms other than  $\rho$  and  $\lambda$ , i.e., clock  $t_{\text{phase}} \in X$  and dense counter  $F_p \in V_{\text{dc}}$ . Clock  $t_{\text{phase}} \in X$  measures the  
 610 total amount of time the automaton spends in location  $h_{\langle \text{op} \rangle}$  and dense counter  $F_p \in V_{\text{dc}}$  is the residual value  
 611 of fatigue at the end of the previous fatigue/recovery cycle, realized by a different  $\langle \text{op} \rangle$ -pub<sub>h</sub> instance. Both  
 612 are updated when a new fatigue/recovery cycle begins, i.e., every time the SHA modeling the human enters  
 613 an instance of  $\langle \text{op} \rangle$ -pub<sub>h</sub> and  $\xi_{\text{rand}, \langle \text{op} \rangle}$  is carried out: clock  $t_{\text{phase}}$  is reset and variable  $F_p$  is updated with  $F$ .

$$\dot{F} = f_{\text{ftg}}(t_{\text{phase}}, \lambda) = F_p \lambda e^{-\lambda t_{\text{phase}}} \quad (4)$$

$$\dot{F} = f_{\text{rec}}(t_{\text{phase}}, \rho) = -F_p \rho e^{-\rho t_{\text{phase}}} \quad (5)$$

614 Dense counter  $f$ , on the other hand, is not associated with a flow in location  $h_{\langle \text{op} \rangle}$ , because it models the  
 615 digital equivalent of the physical attribute  $F$ . For this reason, the temporal evolution of  $f$  is calculated  
 616 explicitly via the update  $\xi_{\langle \text{op} \rangle}$ , which computes a new value for  $f$  by applying the update in Eq.6, every  $T_{\text{poll}}$   
 617 time units. The primed version  $f'$  indicates the new value of  $f$  after the computation of the expression, which  
 618 depends on the operating state  $\langle \text{op} \rangle$ . Unlike Eq.2, the equations in Eq.6 model fatigue in a single cycle and  
 619 are expressed in terms of the amount of time elapsed from the beginning of the current fatigue/recovery  
 620 cycle. Conversely, Eq.2 depends on the absolute time  $t$  and instant  $\mathbf{t}_{i-1}$ , the latter indicating the end of  
 621 the cycle that precedes the current one. Hence, if  $\tau$  is the amount of time elapsed from the beginning of a  
 622 cycle, Eq.2 can be rewritten in terms of  $\tau$  by applying the identity  $t - \mathbf{t}_{i-1} = \tau$ , and fatigue after  $\tau$  time  
 623 units from the beginning of the current fatigue/recovery cycle is  $\bar{F}(\tau) = F(\mathbf{t}_{i-1} + \tau)$ . For  $\tau = 0$ , fatigue  
 624  $\bar{F}(0)$  is equal to the residual value  $F(\mathbf{t}_{i-1})$ , which is  $F_p$ . At the end of the  $(k+1)$ -th sensor refresh, lasting  
 625  $T_{\text{poll}}$  time units each, the fatigue is  $F(kT_{\text{poll}} + T_{\text{poll}})$ . The final expressions are obtained by considering that,  
 626 before computing  $\xi_{\langle \text{op} \rangle}$ ,  $f$  amounts to  $F_p e^{-\rho k T_{\text{poll}}}$ , in case of recovery, and to  $1 - (1 - F_p) e^{-\lambda k T_{\text{poll}}}$  otherwise  
 627 (i.e., the fatigue after  $k$  refresh cycles).  
 628

$$f' = \bar{F}(kT_{\text{poll}} + T_{\text{poll}}) = \begin{cases} 1 - (1 - F_p) e^{-\lambda(kT_{\text{poll}} + T_{\text{poll}})} = 1 - (1 - F_p) e^{-\lambda k T_{\text{poll}}} e^{-\lambda T_{\text{poll}}} = 1 - (1 - f) e^{-\lambda T_{\text{poll}}} & \text{(fatigue)} \\ F_p e^{-\rho(kT_{\text{poll}} + T_{\text{poll}})} = F_p e^{-\rho k T_{\text{poll}}} e^{-\rho T_{\text{poll}}} = f e^{-\rho T_{\text{poll}}} & \text{(recovery)} \end{cases} \quad (6)$$

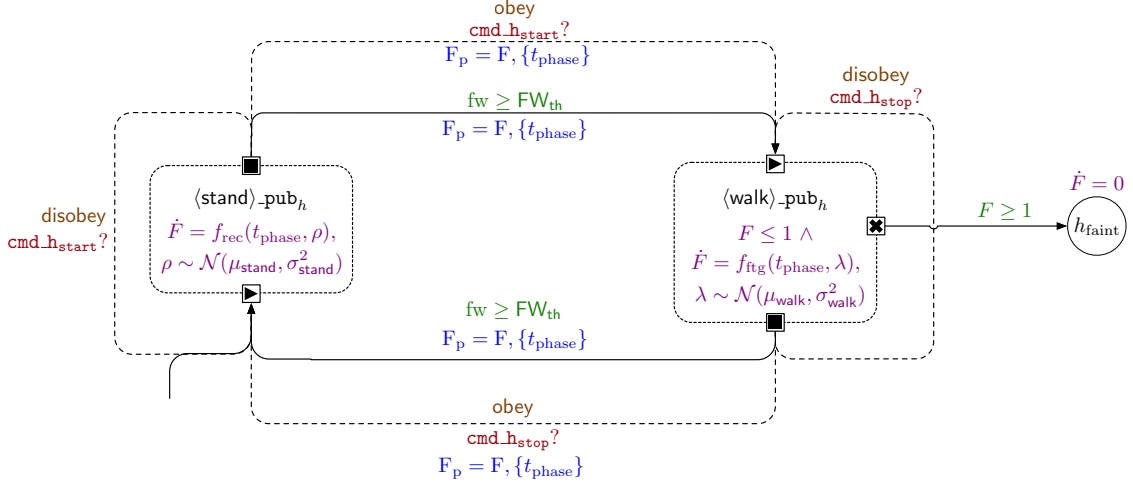


Figure 7: SHA modeling human behavior when adhering to the HumanFollower pattern. Color-coding is the same as Fig. 3.

629 Compared to [13, 14], we extend SHA modeling humans by introducing  $\mathcal{D}(h_{\langle \text{op} \rangle})$ ,  $\xi_{\langle \text{op} \rangle}$ , and  $\xi_{\text{rand}, \langle \text{op} \rangle}$  in all  
630  $\langle \text{op} \rangle_{\text{pub}_h}$  instances. Enriching the SHA with these features strengthens the results obtained with SMC as  
631 they account not only for the uncertainty due to human autonomy, but also for the natural variability of  
632 the fatigue phenomenon. The extension, therefore, leads to more reliable estimations of the fatigue levels  
633 reached by subjects involved in the scenario, including an estimation of their variability ranges.

634 In the following, we present the individual SHA modeling the three interaction patterns, all featuring  
635 multiple instances of the hereby presented  $\langle \text{op} \rangle_{\text{pub}_h}$  pattern.

#### 636 4.1.1. Human Follower

637 An instance of the SHA modeling the human follower pattern is generated for each service specified  
638 through the DSL with `ptrn = HumanFollower`. The SHA, hereinafter referred to as  $\mathcal{A}_{\text{hf}}$  and shown in  
639 Fig. 7, features two instances of the  $\langle \text{op} \rangle_{\text{pub}_h}$  pattern: one capturing the recovery phase while standing  
640 ( $\langle \text{stand} \rangle_{\text{pub}_h}$ ) and one for the fatigue phase while walking ( $\langle \text{walk} \rangle_{\text{pub}_h}$ ). Fatigue decreases while resting  
641 (in  $\langle \text{stand} \rangle_{\text{pub}_h}$ ) and increases while walking (while in  $\langle \text{walk} \rangle_{\text{pub}_h}$ ). Therefore,  $\mathcal{F}(h_{\langle \text{stand} \rangle})$  equals  $f_{\text{rec}}(t, \rho)$   
642 (see Eq.5) and  $\mathcal{F}(h_{\langle \text{walk} \rangle})$  equals  $f_{\text{ftg}}(t, \lambda)$  (see Eq.4). Values  $\rho$  and  $\lambda$  are realizations of  $\mathcal{N}(\mu_{\text{stand}}, \sigma_{\text{stand}}^2)$  and  
643  $\mathcal{N}(\mu_{\text{walk}}, \sigma_{\text{walk}}^2)$ , respectively. Table 1 shows the internal updates,  $\xi_{\text{stand}}$  and  $\xi_{\text{walk}}$ , respectively, later described  
644 in detail.  $\mathcal{A}_{\text{hf}}$  also features a deadlock location  $h_{\text{faint}}$  capturing the case in which the human reaches full  
645 exhaustion causing the failure of the mission. If the mission fails because the human has reached location  
646  $h_{\langle \text{faint} \rangle}$ , modeling the evolution of fatigue is no longer relevant. Therefore, location  $h_{\langle \text{faint} \rangle}$  is endowed with  
647 flow condition  $\dot{F} = 0$ .

648 While walking (thus, while in location  $h_{\langle \text{walk} \rangle}$ ), the SHA periodically updates variables  $h_{\text{pos}_x}$  and  $h_{\text{pos}_y}$ .  
649 As described in Section 3.1, we assume that humans walk at average speed  $v \in K$ . Dense counter  $h_\gamma$   
650 captures the human's orientation with respect to the  $x$ -axis. We assume that the human can rotate instantly  
651 while following their trajectory, and, thus, no location is necessary to capture the delay caused by rotation.  
652 Variable  $h_\gamma$  is periodically updated while walking through function `upd_orientation()`, which computes the  
653 new orientation required (primed dense counter  $h'_\gamma$ ) to head towards the following point of the trajectory.  
654 Therefore, every  $T_{\text{poll}}$  time instants, the  $x$ - $y$  coordinates increase by  $vT_{\text{poll}} \cos(h_\gamma)$  along the  $x$ -axis and  
655  $vT_{\text{poll}} \sin(h_\gamma)$  along the  $y$ -axis. While standing (in  $h_{\langle \text{stand} \rangle}$ ), as per Table 1, the human does not move, thus  
656 the values of  $h_{\text{pos}_x}$  and  $h_{\text{pos}_y}$  do not change. The mechanism capturing human free will and the corresponding  
657 dense counter  $fw$  is presented later in this section. The periodic sensor refresh mechanism does not apply  
658 to location  $h_{\text{faint}}$  (which is not, thus, part of an  $\langle \text{op} \rangle_{\text{pub}_h}$  instance) since, once the mission has failed, the  
659 orchestrator no longer requires up-to-date sensor measurements.

660 The switch between  $h_{\langle \text{stand} \rangle}$  and  $h_{\langle \text{walk} \rangle}$  (and viceversa) is controllable and triggered by events through

Table 1: Updates for the SHA modeling the HumanFollower and HumanLeader patterns.

Symbol	Updates	Description
$\xi_{\text{stand}}$	$f' = f e^{\rho T_{\text{poll}}};$ $h'_\gamma = h_\gamma;$ $h'_{\text{pos}_x} = h_{\text{pos}_x};$ $h'_{\text{pos}_y} = h_{\text{pos}_y};$ $\text{fw}' = \text{roll\_dice}();$	Resting phase
$\xi_{\text{walk}}$	$f' = 1 - (1 - f) e^{-\lambda T_{\text{poll}}};$ $h'_\gamma = \text{upd\_orientation}();$ $h'_{\text{pos}_x} = h_{\text{pos}_x} + v T_{\text{poll}} \cos(h_\gamma);$ $h'_{\text{pos}_y} = h_{\text{pos}_y} + v T_{\text{poll}} \sin(h_\gamma);$ $\text{fw}' = \text{roll\_dice}();$	Fatiguing phase

661 channels `cmd_h_start` and `cmd_h_stop`. The orchestrator sends to the SHA modeling the human events through  
662 these channels when it detects that the interaction between the human and the robot must start (`cmd_h_start`)  
663 or stop (`cmd_h_stop`). Upon switching between  $h_{\langle \text{stand} \rangle}$  and  $h_{\langle \text{walk} \rangle}$ , the SHA updates the value of variable  $F_p$   
664 (see the updates on entering `\langle \text{op} \rangle_{\text{pub}_h}` instances). To capture the unpredictability of human behavior, the  
665 edges between  $h_{\langle \text{stand} \rangle}$  and  $h_{\langle \text{walk} \rangle}$  and back have specific features modeling human free will. In the literature,  
666 there exist several proposals on how to model the free will phenomenon [33]. We exploit the results on the free  
667 will phenomenon and randomness in [34] to model human *haphazard* choices probabilistically. Specifically,  
668 as in Fig. 7, the controllable edges between  $h_{\langle \text{stand} \rangle}$  and  $h_{\langle \text{walk} \rangle}$  and the two self-loops complementing them  
669 are associated with a probability distribution such that  $\text{obey} + \text{disobey} = 1$ , where  $\text{obey}, \text{disobey} \in K$  are two  
670 constants. The values are the probabilities with which, when the orchestrator fires an event over `cmd_h_start`  
671 (resp., `cmd_h_stop`), the human abides by it and switches to  $h_{\langle \text{walk} \rangle}$  (resp.,  $h_{\langle \text{stand} \rangle}$ ) or ignores it and stays in  
672 the same location. The specific value of  $\text{obey}$  derives from attribute  $\mathbf{p}_{\text{fw}}$  of class `Human` introduced in Section  
673 3 (then,  $\text{disobey} = 1 - \text{obey}$  holds). Value `disabled` for attribute  $\mathbf{p}_{\text{fw}}$  implies  $\text{obey} = 1$  (hence,  $\text{disobey} = 0$ ).

674 Manifestations of human free will do not exclusively occur concomitantly with the orchestrator’s instruc-  
675 tions. As shown in Fig. 7, two additional uncontrollable edges connect  $h_{\langle \text{stand} \rangle}$  and  $h_{\langle \text{walk} \rangle}$ . These edges  
676 capture the possibility that humans *may* decide to start or stop walking haphazardly at any time during the  
677 execution of the mission. Therefore, these edges are not associated with any event occurring in the system,  
678 but they only depend on the value of edge condition  $\text{fw} \geq \text{FW}_{\text{th}}$ , where  $\text{FW}_{\text{th}} \in K$  is a constant. As per  
679 Table 1, the value of dense counter  $\text{fw}$  is updated every  $T_{\text{poll}}$  time units through function `roll_dice()`, which  
680 generates a random value from  $[0, \text{FW}_{\text{max}}]$  where  $\text{FW}_{\text{max}} \in K$  is a constant. The uncontrollable edge fires if,  
681 and only if, the generated value of  $\text{fw}$  is greater or equal than constant  $\text{FW}_{\text{th}} \leq \text{FW}_{\text{max}}$ . Parameters  $\text{FW}_{\text{th}}$   
682 and  $\text{FW}_{\text{max}}$  derive from attribute  $\mathbf{p}_{\text{fw}}$  and determine the frequency of human haphazard actions.

#### 683 4.1.2. Human Leader

684 The SHA modeling the leader pattern, shown in Fig. 8, shares most features with the model described  
685 in Section 4.1.1. Locations  $h_{\langle \text{stand} \rangle}$  and  $h_{\langle \text{walk} \rangle}$  (within `\langle \text{stand} \rangle_{\text{pub}_h}` and `\langle \text{walk} \rangle_{\text{pub}_h}`) capture the human  
686 resting and walking constraining real-valued variable  $F$  through the flow conditions in Eq.5 and Eq.4. While  
687 in these locations, sensor readings are periodically modified by updates  $\xi_{\text{stand}}$  and  $\xi_{\text{walk}}$  in Table 1. When  
688 fatigue exceeds the maximum threshold, the SHA switches to deadlock location  $h_{\langle \text{faint} \rangle}$ .

689 The distinguishing feature of this pattern is that the switch from  $h_{\langle \text{stand} \rangle}$  to  $h_{\langle \text{walk} \rangle}$  is purely based on the  
690 free will mechanism and not on orchestrator’s instructions. As a matter of fact, the leader *autonomously*  
691 decides when to start the action. Therefore, the edge to  $h_{\langle \text{walk} \rangle}$  is not tied to any event fired through any  
692 channel. Dense counter  $\text{fw}$  appearing in the edge condition is periodically randomly updated as described  
693 in Section 4.1.1. On the other hand, while the leader is free to also stop walking at any time irrespective  
694 of the robot’s decisions (through the solid edge from  $h_{\langle \text{walk} \rangle}$  to  $h_{\langle \text{stand} \rangle}$ ), the orchestrator may exceptionally  
695 instruct the human to stop walking through channel `cmd_h_stop` when their fatigue reaches an alarming value.  
696 As with all other orchestrator commands, the edges triggered by such events are probabilistic and depend on

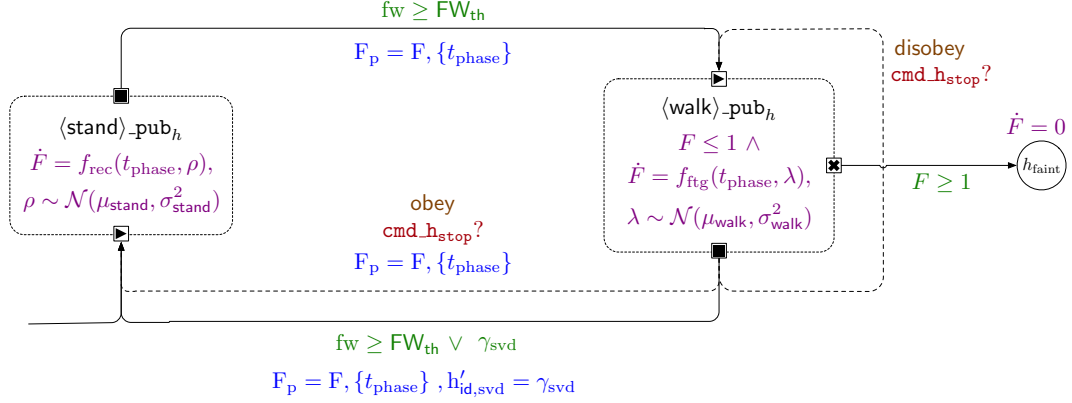


Figure 8: SHA modeling the HumanLeader pattern, color-coded as in Fig. 3.

697 probabilities **obey** and **disobey** (see Fig. 8) governing whether the human abides by the instruction or ignores  
 698 it and stays in the same location.

699 Finally, unlike in the follower pattern, the leader marks the end of the service by updating the Boolean  
 700 dense counter  $h_{id, served} \in V_{dc}$ , which is also used by the orchestrator to determine whether the robot may  
 701 move on to serve the following human or stop if the mission is complete. The condition that determines  
 702 whether the service is complete is indicated as  $\gamma_{svd}$  and corresponds to Formula 7 (see also Fig. 8). The  
 703 service is considered complete if both the human and the robot are within a specific range of the destination,  
 704 corresponding to attribute **target** of class **Service** in Fig. 4. Dense counters  $r_{pos_x}$  and  $r_{pos_y}$  represent the  
 705 Cartesian coordinates of the robot within the layout [15].

$$\sqrt{(h_{pos_x} - \text{target.x})^2 + (h_{pos_y} - \text{target.y})^2} \leq vT_{poll} \wedge \sqrt{(h_{pos_x} - r_{pos_x})^2 + (h_{pos_y} - r_{pos_y})^2} \leq vT_{poll} \quad (7)$$

706 As per Fig. 8, when condition  $\gamma_{svd}$  holds and the edge from  $h_{(walk)}$  to  $h_{(stand)}$  is taken, dense counter  $h_{id, served}$   
 707 is set to 1 (Boolean values are encoded by 0 and 1). If the edge is taken because  $fw \geq FW_{th}$  holds,  $h_{id, served}$   
 708 is set with the value of  $\gamma_{svd}$  (that, possibly, can be 0 if the service has not been completed yet).

#### 709 4.1.3. Human Recipient

710 The recipient pattern captures a human needing the robot to fetch an object and deliver it back to  
 711 their current location. While the robot moves to the object's physical location (i.e., attribute **target** of the  
 712 corresponding **Service**) and travels back, the human is free to move around. Therefore, the SHA modeling  
 713 human behavior for this pattern (shown in Fig. 9) features three operational states, corresponding to as many  
 714 instances of the  $\langle op \rangle\_pub_h$  pattern. Instance  $\langle stand \rangle\_pub_h$  captures the human standing still, as described in  
 715 Section 4.1.1 and Section 4.1.2. Similarly,  $\langle walk \rangle\_pub_h$  captures the human walking out of free will while  
 716 waiting for the robot. Additionally, the recipient pattern features location  $h_{(exe)}$  (within pattern  $\langle exe \rangle\_pub_h$ )  
 717 representing that the robot has reached the human while carrying the object and the human has to collect  
 718 it. During the handover, neither the robot nor the human can move, thus  $\mathcal{F}(h_{(exe)})$  equals  $f_{rec}$ . Ordinary  
 719 location  $h_{(faint)}$  captures the human having reached the maximum fatigue level and, as in previously presented  
 720 patterns, it is endowed with flow condition  $\dot{F} = 0$ .

721 While the robot is busy fetching the object, the human can autonomously decide to move at any time.  
 722 Therefore, the edges from  $h_{(stand)}$  to  $h_{(walk)}$  and back depend on dense counter  $fw$ , which is periodically  
 723 updated as described in Section 4.1.1. The orchestrator starts the handover when the human is ready to  
 724 deliver the object to the robot, by firing an event through channel  $cmd\_h\_start$ . In this case, whether the human  
 725 is walking (thus, in  $\langle walk \rangle\_pub_h$ ) or idle (in  $\langle stand \rangle\_pub_h$ ), they receive the instruction through channel  
 726  $cmd\_h\_start$  to switch to  $\langle exe \rangle\_pub_h$  for the synchronization phase. As with the previous SHA modeling human  
 727 behavior, there is a certain probability that the human ignores the orchestrator's commands as dictated by  
 728 weights **obey** and **disobey**.



754  $\mathcal{A}_r$  has four ordinary non-committed locations, capturing the robot's behavior while: 1. idle (location  
755  $r_{\text{idle}}$ , also corresponding to the initial location); 2. accelerating (location  $r_{\text{start}}$ ); 3. moving at maximum speed  
756 (location  $r_{\text{mov}}$ ); 4. turning (location  $r_{\text{turn}}$ ); 5. decelerating (location  $r_{\text{stop}}$ ). In  $\mathcal{A}_r$ , the robot periodically  
757 shares the updated position values while in  $r_{\text{start}}$ ,  $r_{\text{mov}}$ , and  $r_{\text{stop}}$ . The robot's coordinates within the floor  
758 layout are modeled by two dense-counter variables  $r_{\text{pos}_x}$  and  $r_{\text{pos}_y}$ , which are periodically updated every  
759  $T_{\text{poll}}$  time instants. Interested readers find the graphical representation and detailed description of  $\mathcal{A}_r$  in [15,  
760 Section II.C].

#### 761 4.2.2. Battery Model

762 Mobile robots are typically powered by a lithium battery, which undergoes *charging* and *discharging*  
763 cycles. Therefore, SHA  $\mathcal{A}_b$  modeling the robot's battery, which is presented in this section and shown in Fig.  
764 10, features two ordinary non-committed locations  $b_{\text{dis}}$  and  $b_{\text{rech}}$  corresponding to the discharge and recharge  
765 cycles, respectively, plus a deadlock location  $b_{\text{dead}}$  capturing the case in which the battery is fully discharged.

766 The main physical attribute for a battery is its voltage (representing the charge level), which is modeled  
767 by real-valued variable  $Q$ . Variable  $Q$  is initialized with the initial voltage value  $C_0 \in K$ , i.e., an attribute  
768 of class **Battery** introduced in Section 3. Similarly to fatigue  $F$  in SHA modeling humans, the temporal  
769 dynamics of  $Q$  is determined by flow conditions in  $\mathcal{F}(b_{\text{dis}})$  and  $\mathcal{F}(b_{\text{rech}})$ , whose integral is shown in Eq.8  
770 and Eq.9, respectively. Compared to [13, 14], flow conditions have been refined to match the behavior of  
771 lithium batteries for real robotic devices. As a matter of fact, the entire discharge cycle (from 100% of the  
772 voltage capacity to 0%) can be approximated by an exponential curve [36]. Nevertheless, the real device  
773 is not operational when the voltage drops below a certain threshold (which can vary depending on the  
774 specific battery type and device it is powering), i.e., when the level of charge is not sufficient to power the  
775 wheel motors. Letting the battery pack discharge to very low levels (close to 0%) may actually permanently  
776 damage it [37]. Therefore, compared to [13, 14], we identified equations governing the evolution of variable  
777  $Q$  (shown in Eq.8 and Eq.9) by fitting the discharge/charge curve when the robotic device is operative and  
778 can carry out the assigned mission. A cubic function showed a high fit to the real dynamics. Parameters  
779  $d_{0,1,2,3}, r_{0,1,2,3} \in K$  determining the discharge and recharge curves are fitted based on sensor measurements  
780 collected during charge/discharge cycles of the same robotic device. Parameter  $d_0$  is always set to  $C_0$  at the  
781 beginning of the scenario.

$$782 \quad Q(t) = -d_3t^3 - 2d_2t^2 - d_1t + d_0 \quad (8)$$

$$783 \quad Q(t) = r_3t^3 + 2r_2t^2 + r_1t + r_0 \quad (9)$$

784 The edges from  $b_{\text{dis}}$  to  $b_{\text{rech}}$  and viceversa both model controllable switches, triggered when the orchestrator  
785 fires an event through channels `cmd_b_start` and `cmd_b_stop` instructing the robot to start or stop recharging.  
786 On the other hand, the switch from  $b_{\text{dis}}$  to  $b_{\text{dead}}$  is uncontrollable as it occurs when  $Q = C_{\text{fail}}$  holds, due to  
787 invariant  $Q \geq C_{\text{fail}}$  on  $b_{\text{dis}}$  and condition  $Q \leq C_{\text{fail}}$  on the edge to  $b_{\text{dead}}$ , where  $C_{\text{fail}} \in K$  is an attribute of  
788 class **Battery** (see Section 3). The edge from  $b_{\text{dis}}$  to  $b_{\text{rech}}$  is also constrained by  $Q > C_{\text{fail}}$ , since the automaton  
789 must enter deadlock location  $b_{\text{dead}}$  when  $Q = C_{\text{fail}}$  holds. The edges from  $b_{\text{dis}}$  to  $b_{\text{rech}}$ , and viceversa, are  
790 equipped with updates that initialize  $d_0$  and  $r_0$  from Eq.8 and Eq.9 with the residual charge value from the  
791 previous cycle (i.e., the value of dense counter  $b_{\text{chg}}$ ) and reset clock  $t_{\text{phase}}$ .

792 The battery model features dense counter  $b_{\text{chg}}$ , representing the digital counterpart of  $Q$ , and a modeling  
793 pattern to periodically publish the latest charge measurement governed by clock  $t_{\text{upd}}$  (corresponding to the  
794 `<op>_pub(id)` pattern presented in [15, Section IV.2]). In  $\mathcal{A}_b$ , this occurs while in  $b_{\text{dis}}$  and  $b_{\text{rech}}$  by switching  
795 to committed locations  $b_{\text{pub}_d}$  and  $b_{\text{pub}_r}$ , respectively, when  $t_{\text{upd}} = T_{\text{poll}}$  holds. Upon these switches, clock  
796  $t_{\text{upd}}$  is reset, to begin a new sensor refresh, and dense counter  $b_{\text{chg}}$  is updated through updates  $\xi_{\text{dis}}$  and  $\xi_{\text{rech}}$ .  
797 The new value of the battery charge depends on how many cycles lasting  $T_{\text{poll}}$  time units have been executed  
798 so far, hence how many measurements have been collected. For this reason, automaton  $\mathcal{A}_b$  features a dense  
799 counter  $k \in V_{\text{dc}}$  that keeps track of the number of readings that have been done since the beginning of the  
800 scenario. Updates  $\xi_{\text{dis}}$  and  $\xi_{\text{rech}}$  compute the battery charge at the  $k$ -th refresh cycle  $Q(kT_{\text{poll}})$ . They are  
801 obtained by expanding Eq.9 and Eq.8 when  $t$  is equal to  $(k-1)T_{\text{poll}} + T_{\text{poll}}$ . Unlike the updates in automata  
modeling humans, dependency on index  $k$  cannot be removed in the equations defining  $b'_{\text{chg}}$ . At every sensor




Figure 10: SHA  $\mathcal{A}_b$  modeling the robot's battery behavior. Color-coding is the same as in Fig. 3.

Table 2: Updates for the  $\mathcal{A}_b$  SHA modeling the robot's battery.

Symbol	Updates	Description
$\xi_{dis}$	$b'_{chg} = b_{chg} - T_{poll}((d_3 + 6d_3k)T_{poll}^2 + (d_2 + 2d_2k)T_{poll} + d_1); k = k + 1;$	Discharge phase
$\xi_{rec}$	$b'_{chg} = b_{chg} + T_{poll}((r_3 + 6r_3k)T_{poll}^2 + (r_2 + 2r_2k)T_{poll} + r_1); k = k + 1;$	Recharge phase

802 refresh, dense counter  $k$  is incremented. Updates are shown in Table 2 (where  $b'_{chg}$  is the new value of  $b_{chg}$   
803 after the update). The updated value of  $b_{chg}$  is then published by firing an event through channel  $pub\_b_{ch}$ .

#### 804 4.3. Orchestrator Model

805 The orchestrator controls the robot's behavior based on the current state of the system to drive the  
806 mission to success. As described in previous sections, the humans, the robot, and its battery share sensor  
807 readings with the orchestrator, which checks these values against given policies to determine whether a  
808 certain event has to be fired. Specifically, the orchestrator is fully in control of the mobile robot's behavior  
809 (i.e., it issues every instruction to start or stop moving), while it issues *suggestions* for the human, e.g., to  
810 stop moving when they reach an alarming value of fatigue, which might be dismissed due to human free will.  
811 The degree of intrusiveness of how such suggestions are issued to the subjects must be tailored to the specific  
812 scenario and the involved subjects' demands; however, given the high-level perspective of the framework,  
813 this aspect is currently out-of-scope. An abstract representation of the orchestrator SHA is shown in Fig.  
814 11. The orchestrator operational states (the dashed boxes in Fig. 11) are modeled as submachines. All  
815 the edges connecting them are defined for events of the form  $c!$ , where  $c$  is a channel of the network, as  
816 the orchestrator proactively triggers suitable actions to govern the evolution of the entire scenario. The  
817 orchestrator operational states, i.e., the submachines in it, are described in detail in the following:

- 818 1. the orchestrator is in  $r_{idle}$  when, given the system's state, no action can start and, thus, the robot is  
819 waiting;
- 820 2.  $r_{rech}$  orchestrates the robot's behavior when it has to move to the recharge station and recharge;
- 821 3.  $r_{lead}$  controls the start and the end of the movement when, based on the interaction pattern characterizing  
822 the service underway, the robot *leads* the action (i.e., for the HumanFollower and HumanRecipient  
823 patterns);
- 824 4.  $h_{lead}$  controls the dual case, in which the movement is initiated by the human (i.e., the HumanLeader  
825 pattern);
- 826 5.  $r_{sync}$  controls the robot during the HumanCompetitor pattern [16];



Table 3: Orchestrator start and stop conditions ( $\gamma_{\text{start}}$  and  $\gamma_{\text{stop}}$ , respectively) for each submachine (sub.m.) of the orchestrator. Edge condition  $\gamma$  characterizing a submachine  $x$  is indicated with notation  $x.\gamma$  (e.g.,  $r_{\text{rech}}.\gamma_{\text{start}}$ ).

Sub.m.	Start condition ( $\gamma_{\text{start}}$ )	Stop condition ( $\gamma_{\text{stop}}$ )
$r_{\text{lead}}$	$h_{\text{pattern}} \in \{\text{follower, recipient}\} \wedge$ $\neg h_{\text{served}} \wedge f \leq F_{\text{restart}} \wedge b_{\text{chg}} \geq C_{\text{rech}} \wedge$ $\text{dist}(r_{\text{pos}}, h_{\text{pos}}) \leq D_{\text{restart}}$	$f \geq F_{\text{stop}} \vee b_{\text{chg}} \leq C_{\text{rech}} \vee$ $\left( h_{\text{served}} \wedge \bigvee_{i=1}^{N_h} \neg h_{i,\text{served}} \right) \vee \text{dist}(r_{\text{pos}}, h_{\text{pos}}) \geq D_{\text{stop}}$
$h_{\text{lead}}$	$h_{\text{pattern}} \in \{\text{leader}\} \wedge \neg h_{\text{served}} \wedge$ $b_{\text{chg}} \geq C_{\text{rech}} \wedge h'_{\text{pos}} \neq h_{\text{pos}}$	$f \geq F_{\text{stop}} \vee b_{\text{chg}} \leq C_{\text{low}} \vee$ $\left( h_{\text{served}} \wedge \bigvee_{i=1}^{N_h} \neg h_{i,\text{served}} \right) \vee h'_{\text{pos}} = h_{\text{pos}}$

6.  $hr_{\text{int}}$  controls the robot’s behavior while providing a service that requires precise or close-distance synchronization with the human (i.e., the **HumanApplicant** or **HumanRescuer** patterns) [16].

Submachines in Fig. 11 are endowed with *ports*, intended as in the  $(\text{op})\text{-pub}_h$  pattern. The orchestrator *enters* a submachine through the port marked by symbol “►”, and may *exit* through the ports marked by symbols “■”, “×”, and “✓”, respectively, indicating whether the action has ended (or is momentarily suspended), the mission has ended with failure or with success. Ports highlight the transitions entering and leaving each submachine, constrained by conditions  $\gamma_{\text{start}}$ ,  $\gamma_{\text{stop}}$ ,  $\gamma_{\text{fail}}$ , and  $\gamma_{\text{scs}}$ , each associated with a component-specific formula. The orchestrator enters a submachine when the corresponding  $\gamma_{\text{start}}$  condition is true. If either of  $\gamma_{\text{stop}}$ ,  $\gamma_{\text{fail}}$ , or  $\gamma_{\text{scs}}$  holds, the orchestrator exits the submachine. Locations  $o_{\text{fail}}$  and  $o_{\text{scs}}$  of Fig. 11 correspond to the end of the mission with failure or success, respectively, and are reached when either  $\gamma_{\text{fail}}$  (see Formula 10) or  $\gamma_{\text{scs}}$  (see Formula 11) holds. Failure occurs if, for at least one of the  $N_h \in K$  subjects in the scenario, human fatigue exceeds 1 (i.e.,  $f_i \geq 1$  holds for some  $i$ ) or the robot’s charge drops to a neighborhood of  $C_{\text{fail}} \in K$  (i.e.,  $|b_{\text{chg}} - C_{\text{fail}}| \leq \epsilon$  holds). The latter condition accounts for small fluctuations of the estimated discharge curve.

$$\left( \bigvee_{i=1}^{N_h} f_i \geq 1 \right) \vee |b_{\text{chg}} - C_{\text{fail}}| \leq \epsilon \quad (10)$$

Location  $o_{\text{scs}}$  is reached when the mission has been successfully completed—i.e., when all humans in the scenario have been served: when a human in the scenario with  $\text{id} = i$  is served, the Boolean dense counter  $h_{i,\text{served}} \in V_{\text{dc}}$  is set to true.

$$\bigwedge_{i=1}^{N_h} h_{i,\text{served}} \quad (11)$$

We recall that the main expression whose value we calculate through SMC is  $\mathbb{P}_M(\diamond_{\leq \tau} \text{scs})$ , where Boolean dense counter  $\text{scs}$  is set to **true** upon entering location  $o_{\text{scs}}$  (thus, when the condition in Formula 11 holds). As per Fig. 11, failure is possible for all submachines. On the other hand, only  $r_{\text{lead}}$ ,  $h_{\text{lead}}$ ,  $r_{\text{sync}}$ , and  $hr_{\text{int}}$  have outgoing transitions towards  $o_{\text{scs}}$ , since recharging the robot does not impact service provision (thus, progress towards mission completion).

Submachines  $r_{\text{idle}}$  and  $r_{\text{rech}}$  are presented in detail in [13], while  $r_{\text{sync}}$  and  $hr_{\text{int}}$  are introduced in [16]. Submachines  $r_{\text{lead}}$  and  $h_{\text{lead}}$  are briefly recapped in the following, as they handle interaction patterns covered in this paper and subject to the experimental validation process in Section 6. Table 3 contains the formulae for the start ( $\gamma_{\text{start}}$ ) and stop ( $\gamma_{\text{stop}}$ ) conditions of these submachines. Since the mission is a sequence of services involving a human agent, the orchestrator uses a dense counter  $\text{curr} \in [1, N_h]$  to store the id of the currently served human. Its value is updated by the orchestrator every time a service ends, and the next one can start. The dense counters  $f$ ,  $h_{\text{pattern}}$ ,  $h_{\text{served}} \in V_{\text{dc}}$  and  $h_{\text{pos}} \in V_{\text{dc}}$  keep track of the fatigue of the currently served human, the required interaction pattern, the completion of the service and the position of the human, respectively ( $h_{\text{pos}}$  is a shorthand representing a pair of coordinates); e.g.,  $f = f_i$  holds if  $i$  is equal to  $\text{curr}$ . Fig. 11 highlights the channels through which the orchestrator fires instructions when entering or

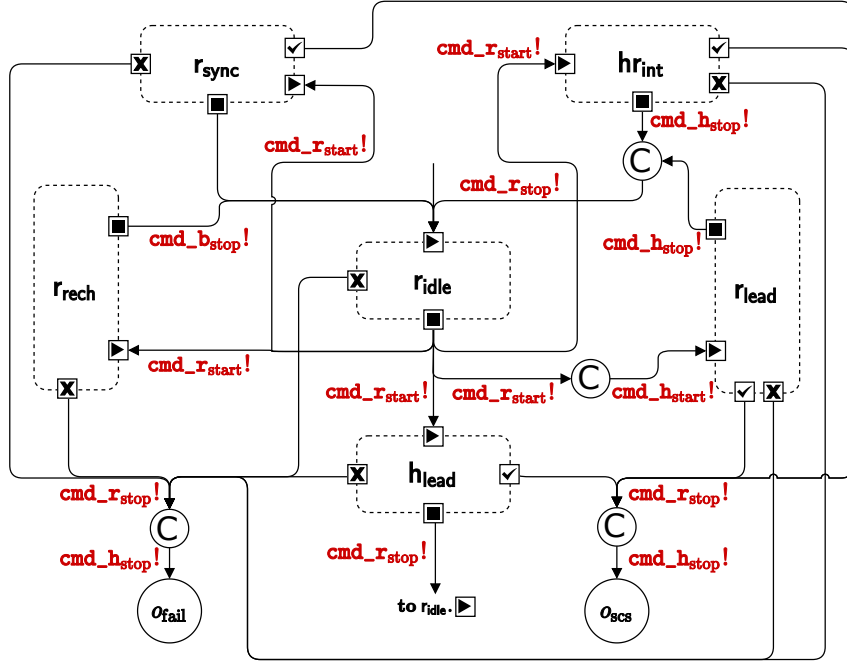


Figure 11: Orchestrator SHA, as seen in [16]. Submachines are represented as dashed boxes, with ports marked by symbols “▶”, “■”, “√”, and “×”.

leaving a submachine. For the sake of clarity, if  $a$  is a submachine, e.g.,  $r_{lead}$ , and  $g$  is the condition on an edge through a port, e.g.,  $\gamma_{start}$ , then we refer to  $g$  by writing  $a.g$ .

The orchestrator enters submachine  $r_{lead}$  to initiate the robot movement when the robot *leads* the action. As per Table 3, the  $r_{lead}.\gamma_{start}$  condition holds for the follower and recipient patterns. The robot begins assisting the currently served human if they are sufficiently close and the service has yet to be completed. Furthermore, for safety purposes, the action can start only if human fatigue is sufficiently low (less than  $F_{restart} \in K$ ) and the robot has sufficient charge (greater than  $C_{rech} \in K$ ). Upon entering  $r_{lead}$ , the orchestrator fires  $cmd\_r\_start$  and  $cmd\_h\_start$  for the robot to start moving and the human to follow. The robot stops moving (events  $cmd\_r\_stop$  and  $cmd\_h\_stop$  fire) if either one of the following conditions holds: (a) human fatigue  $f$  exceeds a maximum tolerable value  $F_{stop} \in K$  (we recall that if the human faints, i.e., their fatigue is 1, the mission fails; hence,  $F_{stop}$  should be lower than 1 to allow the orchestrator to prevent failures caused by fainting); (b) battery charge drops below a value  $C_{rech}$  that calls for recharging; (c) the human has been served, but they were not the last one (if they were the last one, the mission would be complete); (d) the distance between the robot and the human is too large (greater than  $D_{stop} \in K$ ), indicating that the human has stayed behind and needs to get closer to the robot to proceed with the service.

The  $h_{lead}$  submachine controls the robot’s behavior when the human leads the action (i.e., with the *leader* interaction pattern). As per Table 3, the orchestrator enters  $h_{lead}$  (i.e.,  $h_{lead}.\gamma_{start}$  holds) if: (a) the currently assisted human is a leader; (b) the service has not been completed; (c) the robot is sufficiently charged; (d) the human is moving (their current position  $h'_{pos}$  is different from the previous sensor reading  $h_{pos}$ ). Upon entering  $h_{lead}$ ,  $cmd\_r\_start$  is triggered and the robot starts following the human. The orchestrator exits  $h_{lead}$  (i.e.,  $h_{lead}.\gamma_{stop}$  holds) when: (a) the currently served human has reached an excessive fatigue level ( $f \geq F_{stop}$ ); (b) the robot’s battery charge has dropped below the recharge threshold ( $b_{chg} \leq C_{rech}$ ); (c) the human has set themselves as served, but there are other humans to serve in the scenario; (d) the human has stopped moving (their current position  $h'_{pos}$  is the same as the previous sensor reading and the service is yet to be concluded, hence the robots waits for the human to walk again). When  $\gamma_{stop}$  holds, the orchestrator stops the robot through channel  $cmd\_r\_stop$  and instructs the human to stop walking only if they are excessively

885 fatigued. As explained in Section 4.1.2, the human may ignore the orchestrator instruction out of free will.

## 886 5. Scenario Deployment and Reconfiguration

887 In the following, we summarize the main features of the deployment and reconfiguration phases of  
888 the framework (phases **2** and **3** in Fig. 2) to keep the presentation of the validation process in Section 6  
889 self-contained. We refer the interested reader to [15] for the complete description of the deployment approach.

### 890 5.1. Scenario Deployment

891 The goal of the deployment phase is to run the robotic application in a real environment or simulate it  
892 with a realistic physics engine. Also, deployment helps the application designer extract valuable information  
893 about the missions and, possibly, drive a reconfiguration of the scenario, since real executions are available  
894 from the scene. In both cases, executable software is built to run the application. The approach allows for  
895 a *hybrid* deployment environment adhering to the digital-twin paradigm [40] with a real robotic device in  
896 the physical environment interacting with human avatars in the virtual environment. When simulation is  
897 performed, the virtual agents can be controlled by means of specific software components that the simulator  
898 executes to manifest the agents' behavior in the virtual scene. Advanced simulator environments also offer  
899 rich control dashboards that render graphically the virtual 3D scene and allow the user of the simulator  
900 to interact with it through input devices while the scene develops. In our framework, to ensure that the  
901 deployed orchestrator enforces the same policies as in the formal model and that, in case of simulation, the  
902 virtual agents behave correspondingly to their respective SHA, a model-to-code mapping principle converts  
903 every SHA into an executable deployment unit. The latter consists of the executable orchestrator script or  
904 the scripts governing the agents' behavior—either the humans or the robot—within the virtual scene. As  
905 the presence of humans is not always guaranteed and, when human agents are patients in distress, even  
906 discouraged, the application designer performing the simulation directly controls human avatars within the  
907 virtual scene to make them act like real humans in real-world scenarios. The framework allows the designer  
908 to issue commands to the avatars by means of input devices, such as the keyboard, through the scripts which  
909 control them. The human actions modeled with the automata are mapped to keys; keystrokes performed by  
910 the application designer are interpreted by the scripts and then rendered in the scene. To replicate physical  
911 fatigue sensors in the simulated environment and the stochastic behavior of fatigue rates, scripts extract a  
912 random sample from a pool of publicly available electromyography signals and estimate fatigue curves using  
913 the technique described in [38, 39].

914 The deployed orchestrator and the agents communicate over a network of ROS publisher and subscriber  
915 nodes [41] (the “Middleware Layer” in Fig. 2). Each automaton described in Section 4 corresponds to a  
916 deployment unit (e.g.,  $\mathcal{A}_O$  maps to the executable orchestrator and  $\mathcal{A}_R$  to the robot). The firing of an  
917 event through channels in set  $C$  in the formal model (of which Fig. 11 shows an overview) corresponds to  
918 the publication of a message on a ROS topic. More specifically, the deployment unit corresponding to the  
919 “*sender*” SHA (i.e., the one with the edge labeled with  $c!$  with  $c \in C$ ) is the publisher node, whereas the  
920 “*receiver*” SHA (i.e., with the edge labeled with  $c?$ ) is the subscriber node.

921 For each run or simulation of the application, system logs are collected and processed to be examined  
922 by the designer for the reconfiguration phase. Specifically, all data that sensors (either real or simulated)  
923 publish through ROS nodes (i.e., the robot's battery and position values, and all humans' fatigue and position  
924 values) are stored to be examined. The robot carries out the mission by providing services in sequence. The  
925 orchestrator logs relevant events concerning the advancement of each service: when it begins, when it is  
926 completed, when it has to be interrupted and why (either the human is too tired or the battery is too low),  
927 whether the entire mission ends in failure and the source of the failure. Data logged by the orchestrator are  
928 necessary to assess whether the deployed mission has ended with success.

### 929 5.2. Scenario Reconfiguration

930 The reconfiguration phase begins by processing data collected during the deployment to extract metrics  
931 comparable with those calculated at design-time. An example is the success rate observed during the

932 deployment phase to be compared with the value of  $\mathbb{P}_M(\psi)$ , with  $\psi$  expressing the success of the mission.  
933 The designer performs this comparative analysis to assess whether unexpected contingencies emerged at  
934 runtime. If deploying the application highlights unforeseen critical situations, the designer may decide to  
935 **reconfigure** the scenario and iterate all the phases of the workflow in Fig. 2 with the new configuration.

936 Data analysis from deployment and reconfiguration may be necessary as SHA modeling human behavior  
937 have stochastic features that are necessarily an *approximation* of the behavior observable in reality. On the  
938 other hand, the framework targets the *service* level of robotic systems' architectures [35], as it focuses on the  
939 workflow of the mission rather than on aspects related to hardware or structural components: automata  
940 modeling the robot, its battery, and the orchestrator do not thus have stochastic features (i.e., in  $\mathcal{A}_r$ ,  $\mathcal{A}_b$ , and  
941  $\mathcal{A}_o$ , function  $\mathcal{D}(l)$  is undefined for all  $l \in L$  and all edges have probability weight 1). Possible reconfiguration  
942 measures include:

943 **RM1.** Assigning a different **robot** to the task, if the facility has more than one available device. It may  
944 not be feasible for a human subject participating in a service to change their starting position due to  
945 facility policies (e.g., patients necessarily start in the waiting room). On the other hand, two robotic  
946 devices in a fleet may differ either because of their starting position or initial battery charge. In both  
947 cases, this reconfiguration measure can cut down the overall duration of the mission. In the first case,  
948 the robot may require less time to reach the first human to serve. In the second case, the robot may  
949 take less time to recharge, or skip recharging entirely while carrying out the mission.

950 **RM2.** Changing the **order** in which humans are to be served. Note that this is only possible if there are no  
951 logical *dependencies* between the services being swapped. This measure can reduce the overall duration  
952 of the mission if the robot has to cover a smaller distance between services and the maximum level of  
953 fatigue reached by human subjects (thus, impacting their wellbeing), for example if a patient has more  
954 time to rest between two services.


955 **RM3.** Changing the **target** of a pattern, if feasible and compliant with the facility policies. An example is  
956 a patient following the robot directly to the doctor's office without going through the waiting room  
957 first. This reconfiguration measure can reduce both the duration of the mission and the level of fatigue  
958 reached by the human subjects. As a matter of fact, reducing the active time leads to a decrease in the  
959 fatigue endured, since, during a fatigue cycle, time is the only variable in Eq.2.

960 **RM4.** Modifying the orchestrator's **thresholds**. For example, it is possible to reduce the charge threshold  
961 at which the robot is instructed to move to the recharge station or the fatigue level at which the  
962 orchestrator instructs the human to stop walking. The designer must handle the trade-off between  
963 the decrease in mission duration and the increase in probability of failure (for example, due to battery  
964 degradation).


965 The cyclical nature of the framework allows the analyst to modify the scenario and perform multiple  
966 iterations of the analysis until verification results are deemed acceptable. The framework also supports the  
967 designer in terms of which and how many parameters of the scenario require manual specification. Parameters  
968 concerning the robot (i.e., speed and acceleration) and battery behavior (i.e., charge and discharge rates) are  
969 provided by the framework to decrease the manual effort required on the designer's side. Designers manually  
970 specify parameters concerning the specific robotic mission (i.e., how many humans are involved and the  
971 service they request) whose value cannot be known a-priori by the framework. Should the tuning of such  
972 parameters result overly cumbersome for the practitioner, splitting the mission into smaller sequences to be  
973 individually analysed is feasible.

## 974 6. Experimental Validation

975 This section presents the results obtained while validating the presented approach on case studies inspired  
976 by the healthcare setting. The validation process addresses the following questions:



(a) Points of interest within the experimental setting: the entrance and recharge station are in orange, cupboards are in green (CUP1 and CUP2), examination/waiting room entrances are in red (R1a, R1b, and R2), and doctors' offices in blue (OFF1, OFF2, and OFF3).



(b) Representation of the layout abstraction as a set of rectangular areas, with wall sizes (m) and intersection points between areas marked by  $\times$  symbols.


Figure 12: Layout used for the experimental validation.

- 977 **G1.** Is the formal model presented in Section 4 adherent to the physical robotic system and how accurate  
 978 are SMC results?
- 979 **G2.** Is the model-driven framework described in Section 3 and Section 5 practical and useful while developing  
 980 interactive scenarios with multiple subjects and services? More specifically, we evaluate:
- 981 (a) how the DSL supports designers in configuring complex scenarios;
- 982 (b) how the design-time analysis phase provides reliable and valuable insights into the modeled  
 983 scenario;
- 984 (c) how scenarios can be reconfigured to improve key indicators (the probability of success and  
 985 estimated fatigue level of human subjects).

986 Both validation phases have been carried out following the framework workflow in Fig. 2. Firstly, we  
 987 model the scenarios through the DSL described in Section 3. The case studies feature one mobile robot  
 988 providing services (in compliance with the patterns presented in Section 3.1) to one or multiple humans.  
 989 Agents operate within the floor layout represented in Fig. 12, corresponding to the third floor of Building 22  
 990 of Politecnico di Milano. Specifically, Fig. 12a highlights the POIs, while Fig. 12b shows how the real layout  
 991 is abstracted as a set of rectangular areas as described in Section 3.2.1. While the physical layout where  
 992 the robotic device moves is a university building, its areas are repurposed in the simulation environment to  
 993 reproduce a healthcare setting. The layout features a main entrance ENTR where the robot meets patients to  
 994 assist them and two side aisles, each with cupboards containing medical kits (CUP1 and CUP2), two rooms  
 995 serving either as waiting rooms for the patients or examination rooms where doctors administer medications  
 996 (R1a, R1b, and R2), and doctors' offices (OFF1, OFF2, and OFF3).

997 DSL models<sup>5</sup> are automatically converted into a JSON file and finally into an Uppaal model implementing

<sup>5</sup>The DSL sources are available at: <https://github.com/LesLivia/hri.dsl>.



(a) Portion of the simulation scene with the simulated human and the phantom robot replicating the real robot’s behavior. (b) Real TurtleBot3 operating in the physical environment and reacting to the human in the simulation scene.

Figure 13: Hybrid real/simulated deployment environment.

998 the SHA network described in Section 4 modeling the specific scenario.<sup>6</sup> The framework also automatically  
 999 sets up and runs the SMC experiment. For the case studies discussed in this section, we perform SMC  
 1000 through Uppaal v.4.1.26 on a machine running macOS v.10.15.7 with 4 cores and 8GB of RAM.

1001 All case studies are subsequently deployed as described in Section 5.1 [15].<sup>7</sup> We have adopted the digital-  
 1002 twin deployment pattern [42] (see Fig. 13) with a real mobile robot operating in the physical environment  
 1003 (shown in Fig. 13b) and reacting to virtual human subjects in the simulation scene (of which a portion is  
 1004 shown in Fig. 13a). The hybrid deployment environment allows us to verify the adherence of the robotic  
 1005 system’s model and the orchestrator’s efficacy with a real device while also performing several runs with  
 1006 (virtual) subjects exhibiting critical fatigue profiles. Electromyography signals serving as dataset to simulate  
 1007 fatigue curves in the simulation environment are provided by [43]. The mobile device is a TurtleBot3 Waffle  
 1008 Pi.<sup>8</sup> Scenarios are deployed using V-REP v.3.6.2 for the simulation scene, Python v.3.6.9 for the orchestrator  
 1009 script, and ROS Melodic to communicate with virtual agents and the TurtleBot3 [41]. The deployment  
 1010 software tools run on a single machine running Ubuntu v.18.04 with 2 cores and 4GB of RAM.

### 1011 6.1. Formal Model Validation

1012 The purpose of the hereby presented experiments is to assess the accuracy of the formal model presented  
 1013 in Section 4 and of the SMC results (validation goal **G1**). To this end, we perform the design-time analysis  
 1014 as presented in Section 3 on three scenarios, referred to as HF (“Human Follower”), HL (“Human Leader”)  
 1015 and LB (“Low Battery”), all taking place in the experimental setup in Fig. 12 and described in detail in  
 1016 Table 4. The three scenarios are structured to validate the main features of the formal model: human-robot  
 1017 dyadic leader/follower dynamics, human fatigue model, robot’s battery model, and orchestrator policies.  
 1018 Therefore, the three experiments test the formal model with both critical (low battery in LB and high fatigue  
 1019 in HL) and non-critical elements (high battery in HF/HL and low fatigue in HF/LB).

1020 We perform SMC with decreasing values of time bound  $\tau$  to estimate the success probability of the  
 1021 three scenarios (i.e., query with `QueryType P_SCS`). This corresponds to the value of expression  $\mathbb{P}_M(\diamond_{\leq \tau} \text{scs})$ ,  
 1022 where  $M$  is the SHA network modeling each mission. SMC experiments are performed in Uppaal with

<sup>6</sup>The software tool to automatically generate the Uppaal model is available at: [https://github.com/LesLivia/hri\\_designtime](https://github.com/LesLivia/hri_designtime).

<sup>7</sup>The software tool implementing the deployment approach is available at: [https://github.com/LesLivia/hri\\_deployment](https://github.com/LesLivia/hri_deployment)

<sup>8</sup>Full documentation and technical specifications available at: <https://emmanual.robotis.com/docs/en/platform/turtlebot3/overview/>.

Table 4: Scenarios used for the formal model validation phase (abbreviation and detailed description). For each service, we indicate the starting location of the human and the target location as START → TARGET.

SCENARIO	DESCRIPTION	MISSION
HF	The robot (Tbot) <i>leads</i> the human (H1) from OFF1 to CUP1. The robot is sufficiently charged to complete the mission, and the human exhibits a non-critical fatigue profile (Young/Healthy).	H1 Follower OFF1 → CUP1
HL	The robot (Tbot) <i>follows</i> the human (H1) from OFF1 to CUP1. The robot is sufficiently charged to complete the mission, and the human exhibits a critical fatigue profile (Elderly/Sick).	H1 Leader OFF1 → CUP1
LB	The robot (Tbot) <i>leads</i> the human (H1) from OFF1 to CUP1. The robot gets fully discharged during the mission, while the human exhibits a non-critical fatigue profile (Young/Healthy).	H1 Follower OFF1 → CUP1

---

**Algorithm 1** Estimation of the the success probability CI for a set of deployment traces  $\mathcal{DT}$ .

---

**Input:**  $\mathcal{DT}$ ,  $\tau$ ,  $N_h$ ,  $T_{\text{int}}$ ,  $\alpha$

```

1:  $\mathcal{DT}_{\text{scs}} \leftarrow \emptyset$ 
2: for  $dt \in \mathcal{DT}$  do
3:    $\mathcal{SVD}_{dt} \leftarrow \{t | t \in \mathbb{N} \wedge i \in [1, N_h] \wedge h_{i,\text{svd}} \in dt(t)\}$   $\triangleright$  Timestamps corresponding to the end of a service.
4:   if  $|\mathcal{SVD}| = N_h \wedge \max(\mathcal{SVD}) \leq \tau - T_{\text{int}}$  then
5:      $\mathcal{DT}_{\text{scs}} \leftarrow \mathcal{DT}_{\text{scs}} \cup \{dt\}$   $\triangleright$  All humans have been served within  $\tau - T_{\text{int}}$ .
6:   end if
7: end for
8:  $p_l \leftarrow \text{ppf}(\alpha/2, |\mathcal{DT}_{\text{scs}}|, |\mathcal{DT}| - |\mathcal{DT}_{\text{scs}}| + 1)$ 
9:  $p_u \leftarrow \text{ppf}(1 - \alpha/2, |\mathcal{DT}_{\text{scs}}| + 1, |\mathcal{DT}| - |\mathcal{DT}_{\text{scs}}|)$ 
10:  $\epsilon \leftarrow (p_u - p_l)/2$ 
11:  $p \leftarrow p_l + \epsilon$ 
Output:  $p, \epsilon$ 

```

---

1023 default statistical parameters, thus with  $\epsilon = \alpha = 0.05$ . We also estimate the maximum human fatigue  
1024 value (expression  $E_{M,\tau}[\max(f)]$ ), and the robot’s residual charge at the end of the mission (expression  
1025  $E_{M,\tau}[\min(b_{\text{chg}},\%)]$ ), i.e., queries with QueryType E\_FTG and E\_CHG, respectively.

1026 Subsequently, we deploy the three scenarios in the digital-twin setting (see Fig. 13) to collect runtime  
1027 observations, compute the same metrics, and compare the results with those obtained at design-time. To  
1028 this end, we apply a partial replication of SMC (summarized by Algorithm 1) to the traces collected through  
1029 deployment to estimate the success probability range observed at runtime. We refer to the simulation log and  
1030 sensor logs collected during a single run (described in Section 5.1) as **deployment trace**. Given deployment  
1031 trace  $dt$ , we indicate as  $dt(t)$  the set of data (sensor readings and milestones recorded by the orchestrator, if  
1032 any) logged at time  $t \in \mathbb{N}$ . Since the orchestrator records the timestamp at which each human is served, it is  
1033 possible to infer from a deployment trace  $dt$  whether the mission ended successfully. If human  $i \in [1, N_h]$  has  
1034 been served in trace  $dt$ , there exists  $t \in \mathbb{N}$  such that  $h_{i,\text{svd}} \in dt(t)$  holds. We indicate as  $\mathcal{DT}$  the set of all  
1035 deployment traces collected for a given scenario. Set  $\mathcal{SVD}_{dt} = \{t | t \in \mathbb{N} \wedge i \in [1, N_h] \wedge h_{i,\text{svd}} \in dt(t)\}$  at Line  
1036 3 in Algorithm 1 contains the timestamps corresponding to the completion of a service in a specific trace  $dt$ .  
1037 Similar to SMC, given a time-bound  $\tau$  and the set of collected deployment traces  $\mathcal{DT}$ , for each deployment  
1038 trace  $dt \in \mathcal{DT}$  we check whether the mission has ended with success within  $\tau$  (i.e., whether  $\diamond_{\leq \tau} \text{scs}$  holds  
1039 for  $dt$ ). Algorithm 1 checks through the condition on Line 4 whether set  $\mathcal{SVD}_{dt}$  has  $N_h$  elements (i.e., *all*  
1040 humans have been served) and the maximum of  $\mathcal{SVD}_{dt}$  is smaller than  $\tau - T_{\text{int}}$  (i.e., the *last* human to be  
1041 served has been served within the time bound minus the time required by the orchestrator to process the  
1042 information). If condition on Line 4 is verified, trace  $dt$  constitutes a success and is added to set  $\mathcal{DT}_{\text{scs}}$  by

Table 5: Comparison between the results obtained through SMC at design time (DT) and the results obtained by deploying the three model validation scenarios (DEPL). For decreasing values of time bound  $\tau$  ([s]), the table contains the verification time ([min]), the probability CI estimated through Uppaal, the CI observed at runtime, and the runs necessary to compute such estimations. The table also contains the estimated maximum human fatigue values ( $[0 - 1]$ ) and minimum charge levels ([%]). For each metric, configurations leading to the least accurate results are highlighted in grey.

SC.	$\tau$	Ver.Time [min]	Success Probability		Runs		Max. Fatigue		Min. Charge	
			DT	DEPL	DT	DEPL	DT	DEPL	DT	DEPL
HF	75	0.41	0.952 ± 0.05	0.950 ± 0.05	29	110	0.0208 ± 0.004	0.0211 ± 0.007	95.54%	95.75%
	53	3.42	0.859 ± 0.05	0.865 ± 0.06	199	110	0.0177 ± 0.005	0.0179 ± 0.008	96.86%	96.27%
	50	3.23	0.812 ± 0.05	0.800 ± 0.07	250	110	0.0167 ± 0.003	0.0168 ± 0.008	97.05%	96.51%
	40	4.25	0.433 ± 0.05	0.500 ± 0.10	395	110	0.0125 ± 0.004	0.0125 ± 0.008	97.65%	97.11%
	34	2.19	0.239 ± 0.05	0.252 ± 0.08	296	110	0.0112 ± 0.003	0.0114 ± 0.003	98.01%	98.10%
HL	50	0.31	0.952 ± 0.05	0.950 ± 0.05	29	120	0.2015 ± 0.037	0.1939 ± 0.050	73.71%	73.32%
	42	2.72	0.826 ± 0.05	0.840 ± 0.07	236	120	0.1763 ± 0.038	0.1623 ± 0.051	74.19%	74.17%
	38	2.46	0.676 ± 0.05	0.664 ± 0.09	354	120	0.1599 ± 0.014	0.1577 ± 0.025	74.43%	74.38%
	35	3.94	0.409 ± 0.05	0.395 ± 0.09	389	120	0.1545 ± 0.042	0.1561 ± 0.027	74.61%	74.55%
	33	2.13	0.216 ± 0.05	0.208 ± 0.07	277	120	0.1451 ± 0.045	0.1434 ± 0.025	74.74%	74.91%
LB	150	1.02	0.000 ± 0.05	0.000 ± 0.05	29	107	–	–	0.001%	–0.001%

1043 the instruction on Line 5.

1044 Algorithm 1 computes  $\mathbb{P}_{\mathcal{DT}}(\diamond_{\leq \tau} \text{scs})$  in terms of a Bayesian Confidence Interval (CI) of the form  $p \pm \epsilon$   
1045 with confidence level  $1 - \alpha$ . We adopt the Clopper-Pearson approach for binomial distributions to compute  
1046 the CI as it is also exploited by the Uppaal tool. Specifically,  $p_l = p - \epsilon$  can be calculated as the  $\alpha$ -quantile  
1047 of a Beta distribution with parameters `successes` and `trials - successes + 1` (Line 8), while  $p_u = p + \epsilon$  can  
1048 be calculated as the  $\gamma$ -quantile with  $\gamma = 1 - \alpha$  and parameters `successes + 1` and `trials - successes` (Line 9)  
1049 [45].<sup>9</sup> Unlike point estimator `successes/trials`, this procedure also provides an insight into the variability of  
1050 the success rate (i.e., the value of  $\epsilon$ ) and how its value changes as more runs are performed.

1051 The results of the SMC experiments, the time and runs necessary to complete it with the required level of  
1052 confidence, and the fatigue and charge estimations are reported in Table 5 (marked as DT, “Design Time”).  
1053 The success probability ranges estimated for scenarios HF, HL, and LB through Algorithm 1 are reported in  
1054 Table 5 (marked as DEPL, “Deployment”).

1055 Results in Table 5 corroborate the intuition that, for decreasing values of  $\tau$ , the probability of success  
1056 decreases both at design time and during deployment. Experimental results with the largest difference  
1057 between design-time and deployment estimations are highlighted in grey. We select values of parameter  $\tau$  to  
1058 be displayed in Table 5 corresponding to probability ranges in three macro-intervals: *high* success probability  
1059 (i.e., with  $p > 80\%$ ), *average* success probability (i.e., with  $40\% < p < 70\%$ ), and *low* success probability (i.e.,  
1060 with  $p < 25\%$ ). Scenarios HF and HL require 75s and 50s, respectively, to end successfully with probability  
1061 approximately equal to 1, while it drops to approximately 20% when the analysis is bounded to 34s and 33s.  
1062 The variability of the success probability between runs within 34s (for HF) or 33s (for HL) and those requiring  
1063 up to 75s is due to the human stopping haphazardly during the mission as described in Section 4.1, causing  
1064 a delay in the completion of the mission. As shown in Table 5, the configurations with the largest difference  
1065 between design-time and runtime estimations of the success probability (highlighted in grey) are also the  
1066 ones requiring the largest number of traces for the SMC experiment (395 and 389 compared to the 110 and  
1067 120 performed in the physical setting). This is due to how  $\mathbb{P}_M(\diamond_{\leq \tau} \text{scs})$  and  $\mathbb{P}_{\mathcal{DT}}(\diamond_{\leq \tau} \text{scs})$  are calculated: if  
1068 traces that have been generated or collected are consistent with each other (e.g., they are all successful),  
1069 it takes a smaller set of traces to obtain a certain value of  $\epsilon$  than when the number of successes fluctuates.  
1070 Indeed, the estimated success probabilities resulting from the configurations requiring the largest number of  
1071 runs (395 for HF and 389 for HL) are also the closest to 50%. In the worst case, the values of  $p$  estimated at

<sup>9</sup>The Python implementation exploits the `scipy.stats.distributions.beta.ppf` function (referred to as `ppf` in Algorithm 1) from the SciPy library to calculate the required quantiles. Full documentation available at: <https://docs.scipy.org>.






Figure 14: Graph representing the battery voltage ( $Q$  [V]) evolution during a complete TurtleBot3 discharge/recharge cycle (approximately 140min each). Dots represent real voltage sensor readings. The red and green lines are the fitted discharge and recharge curve, respectively. Black dashed lines mark the voltage values corresponding to 100% (about 12.4V) of the charge and 0% (11.0V). The red dashed line marks the time instant where the design-time estimation is the least accurate.

design-time and runtime differ by 6.7% (for HF) and 1.4% (for HL) while this drops to 0.2% in the best case.

Data collected through the three scenarios are also necessary to assess whether the formal model accurately captures the robot’s battery voltage drop (i.e., the battery discharging while the robot is operative). To estimate the expected value of the minimum charge for the same decreasing values of  $\tau$  used for the success probability ranges, we calculate the value of expression  $E_{M,\tau}[\min(b_{\text{chg},\%})]$  in Uppaal (column DT) and the average of minimum values logged for each deployment trace  $E_{DT,\tau}[\min(b_{\text{chg},\%})]$  (column DEPL). Table 5 reports the battery charge percentage estimations at time  $\tau$  for the three scenarios, highlighting, as in previous cases, the configurations leading to the largest estimation error. We recall that the differential equations obtained by deriving Eq.8 and Eq.9 constrain the battery voltage ([V]), whose value in the real system is directly measured through a sensor. The percentages shown in Table 5 are calculated according to Eq.12, where  $b_{\text{chg}}$  represents the dense counter presented in Section 4.2.2 (for column DT) or the value shared by the real sensor (for column DEPL),  $C_{\text{fail}}$  is the lowest voltage value allowed by the device (as presented in Section 4.2.2), and  $C_{\text{full}}$  is the (approximate) voltage value when the battery is fully charged.

$$b_{\text{chg},\%} = \frac{b_{\text{chg}} - C_{\text{fail}}}{C_{\text{full}} - C_{\text{fail}}} \cdot 100 \quad (12)$$

For the three scenarios, we estimate the residual battery charge, assuming that  $C_0$  is set to 99%, 75%, 0.8% for HF, HL, and LB, respectively. The largest estimation error is 0.61% for HF and 0.53% for HL. Unsurprisingly, charge-related estimations are more accurate than for human fatigue, which has a higher degree of variability and is subject to the human’s unforeseeable choices. Nevertheless, the time span required for the scenarios is orders of magnitude shorter than a full battery discharge cycle (approximately 2.5h). Therefore, while these estimations provide insights into how accurate the model is in the range of seconds, we have also assessed its accuracy in the longer run. We have recorded the real battery sensor readings over the course of three full discharge/recharge cycles. This set of data has been used to fit parameters in Eq.8 and Eq.9 governing the evolution of real-valued variable  $Q$  in the SHA. Fig. 14 shows the voltage curves (modeled in SHA through real-valued variable  $Q$ ) resulting from the fitting (red and green lines), compared against the actual sensors readings of a fourth complete discharge/recharge cycle (grey dots). Sensor readings used to fit the curve parameters and those shown in Fig. 14 are different data sets. The largest estimation difference (also highlighted in Fig. 14) is 0.54%, which is comparable to the previously described differences obtained with the three scenarios.

Scenario LB requires a separate analysis. The purpose of this experiment is to assess whether the formal model accurately captures reality in a boundary condition where the robot’s charge is insufficient to complete the mission (as previously mentioned, in this case  $C_0$  is 0.8%). When the mission starts, since the charge level is too low ( $c \leq C_{\text{low}}$  holds), the orchestrator immediately instructs the robot to start moving towards the recharge station, which, however, requires about 3.5min to be reached while the robot has only 2.5min of battery life left. The design-time analysis correctly predicts this outcome as the mission has 0% success

1105 probability within 150s (see Table 5), and all the collected deployment traces end in failure. No fatigue  
 1106 estimation is provided in this case since the human never starts moving. As a matter of fact, the robot needs  
 1107 to recharge as soon as the mission starts, thus the orchestrator immediately enters submachine  $r_{\text{rech}}$  (see Fig.  
 1108 11) without sending any instructions to the human. As per Formula 10, failure occurs when voltage drops  
 1109 sufficiently close to 0%, which is why the estimation reported in Table 5 is not exactly 0% but the estimated  
 1110 probability of failure is still 1. On the other hand, the negative percentage estimated from deployment traces  
 1111 is due to how the real device works. As soon as the detected battery voltage equals 11V, the device will  
 1112 start emitting an acoustic signal to notify the need to recharge, it will beep for a few seconds and then stop  
 1113 sending power to the motors (thus, no motion is possible). From the moment it starts beeping to the moment  
 1114 it stops moving, the voltage drops slightly *below* 11V, leading to the negative percentage (see Eq.12).

1115 Concerning the estimation of the human fatigue, Table 5 reports the maximum value estimated through  
 1116 Uppaal (column DT) and from deployment traces (column DEPL). To estimate the maximum fatigue expected  
 1117 value at design time, we compute the value of expression  $E_{M,\tau}[\max(f)]$  in Uppaal (Table 5, column DT),  
 1118 whose result is a 95% confidence interval of the requested value [17]. The same procedure is applied to the  
 1119 set of deployment traces  $\mathcal{DT}$ . For each deployment trace, we calculate the maximum value of fatigue of the  
 1120 human subject within time bound  $\tau$ . The so-obtained values constitute the set of independent samples, of  
 1121 which we subsequently calculate the 95% confidence level (results are reported in Table 5, column DEPL). In  
 1122 scenario HF, the human has the least critical fatigue profile (Young/Healthy), thus it only reaches a fatigue  
 1123 value of approximately 2%. For this scenario, in the worst case, the largest design-time estimation error  
 1124 (calculated as difference between the average value at design time and observed during deployment) is 1.75%.  
 1125 In HL, the human reaches higher fatigue values (up to approximately 20%). All intervals calculated from  
 1126 deployment traces fall within the range estimated at design-time. In the worst case, also highlighted in Table  
 1127 5, the estimation error is 8.6%. Although design-time estimations pertaining to fatigue are promising, it is  
 1128 important to remark that the simulator scripts governing human behavior during deployment directly result  
 1129 from the model-to-code transformation (presented in [15]) of the SHA described in Section 4.1. These results,  
 1130 therefore, do not constitute a conclusive empirical proof that SHA modeling humans accurately capture  
 1131 reality. Nevertheless, since simulated sensors share their readings with the orchestrator over actual ROS  
 1132 topics, the limited design time-to-deployment errors indicate that modeling patterns dealing with readings  
 1133 update and publishing (i.e., the pattern in Fig. 6 and the `RosPubNode` pattern presented in [15]) are reliable.

1134 Concerning performance and scalability, given the limited complexity of the scenarios analysed in this  
 1135 batch of SMC experiments, verification experiments performed through Uppaal last between 0.31 and 4.25min.  
 1136 For the deployment phase, we have performed a total of 337 real runs (110 for HF, 120 for HL, 107 for LB).  
 1137 By factoring in the time required to perform each run, reset the layout to its starting configuration at the  
 1138 end of each run, and the time required to recharge the robot, this corresponds to approximately 64h of  
 1139 non-stop deployment and runtime data collection. Considering that, in a real healthcare facility, employees  
 1140 have multiple tasks to deal with, robots are not actively deployed 100% of the time, and there are time  
 1141 breaks between shifts, the data collection phase would take a longer time.

## 1142 6.2. Model-Driven Framework Validation

1143 After the analysis on the accuracy of the formal model, we focus on the overall efficacy of the model-driven  
 1144 framework (goal **G2**). To this end, we have analysed 41 real-world scenarios describing service robotic  
 1145 applications extracted from [46, 47, 48, 49, 50, 51]. We remark that, given that service robot deployment is not  
 1146 widespread at the time of writing, there is no structured repository collecting natural language specifications  
 1147 of such scenarios; to the best of the authors' knowledge, the RoboMAX repository (providing 14 of the  
 1148 41 identified scenarios) is the first attempt in this direction [47]. Therefore, the scenario collection phase  
 1149 has been performed manually by surveying related works in the literature and commercial service robots'  
 1150 documentation. Within the set of eligible scenarios, we consider 14 scenarios to fall outside of the scope of our  
 1151 framework. As our work targets service robot applications featuring mobile robots that *interact* with humans,  
 1152 we consider out-of-scope all scenarios featuring robots that operate autonomously (e.g., performing patrolling  
 1153 or automated room cleaning), are teleoperated, or provide information without expecting any reaction on the  
 1154 human side (e.g., periodical medication reminders). As for the 27 scenarios that are within the scope of the  
 1155 framework, we assess that 24 can be modeled through our framework, leading to a coverage percentage of

Table 6: Scenarios used for the framework validation phase (abbreviation, detailed description, and sequence of services). For each service, we indicate the starting location of the human and the target location as START → TARGET.

SCENARIO	DESCRIPTION	MISSION
DPa	The robot (Tbot) serves a patient-doctor pair (P1/D1, respectively). The robot meets the patient by the entrance (ENTR) and <i>leads</i> them to the waiting room (R1b) to wait for the doctor to visit them. The robot <i>follows</i> the doctor to CUP1 where they fetch required tools, and <i>follows</i> them back (carrying the tools) to the examination room (R2) where the patient will receive the treatment. Finally, the robot returns to R1b and <i>escorts</i> the patient to R2 where the doctor is waiting.	P1 Follower ENTR → R1b, D1 Leader R2 → CUP1, D1 Leader CUP1 → R2, P1 Follower Rb1 → R2
DPb	The robot (Tbot) serves a patient-doctor pair (P1/D1, respectively). The robot meets the patient by the entrance (ENTR) and <i>leads</i> them to the waiting room (R1a) to wait for the doctor to visit them. The robot approaches CUP2 to retrieve a required medical kit, and then <i>delivers</i> it to D1 at OFF2. The robot <i>follows</i> the doctor to the examination room (R2) where the patient will receive the treatment. Finally, the robot returns to R1a and <i>escorts</i> the patient to R2 to be treated.	P1 Follower ENTR → R1a, D1 Recipient OFF2 ↔ CUP2, D1 Leader OFF2 → R2, P1 Follower R1a → R2
DPc	The robot (Tbot) serves two patient-doctor pairs (P1 and P2 are patients, D1 and D2 are doctors). The robot meets P1 by the entrance (ENTR) and <i>leads</i> them to the waiting room (R1a), then it performs the same task for P2 <i>leading</i> them from the entrance to R1b. The robot fetches the first required medical kit from CUP1 and <i>delivers</i> it to D1 at OFF1. The robot then serves D2 by <i>following</i> them to CUP2 and back to their office (OFF3) while carrying the kit. Finally, the robot <i>leads</i> P1 to OFF1 and P2 to OFF3 as both doctors are ready to visit them.	P1 Follower ENTR → R1a, P2 Follower ENTR → R1b, D1 Recipient OFF1 ↔ CUP1, D2 Leader OFF3 → CUP2, D2 Leader CUP2 → OFF3, P1 Follower R1a → OFF1, P2 Follower R1b → OFF3

1156 88.8%. The three scenarios that our framework does not cover feature: a) *exoskeletons*, which are considered  
1157 service robots by standard ISO 13482 [52] (thus, they are in-scope with respect to our framework) but  
1158 require different software development practices than mobile robots that our framework targets; b) *cognitive*  
1159 interaction (e.g., comforting children or rehabilitating cognitive skills of patients recovering from strokes),  
1160 whereas our framework targets *physical coordination* between humans and robots.

1161 To address goal **G2**, we have developed three more sophisticated scenarios, referred to as DPa (“Doctor-  
1162 Patient”), DPb, and DPc, collecting the most frequent elements of the 24 real-world scenarios found in the  
1163 literature (fetch-and-delivery tasks, doctor-patients dynamics, patient greeting, and transporting items). The  
1164 three scenarios are described in detail in Table 6. In all three scenarios, the robot has to serve one (in DPa  
1165 and DPb) or two (DPc) pairs of human subjects representing a doctor and a patient. The robot always  
1166 accompanies the patient to the waiting room (R1a, R1b, or R2) first (adhering, thus, to the Follower pattern),  
1167 and then supports the doctor in retrieving the instrumentation needed to treat the patient. Doctors are  
1168 either Leaders (D1 in DPa and DPb, and D2 in DPc) or Recipients (D1 in DPc) depending on whether they  
1169 personally lead the robot to destination or it moves independently and then delivers the resource. For this  
1170 experimental phase, the robot’s charge is always sufficient (the opposite boundary condition has already

Table 7: Results of the DSL2SHA calculation, of the design-time analysis (DT) and of the deployment phase (DEPL) for scenarios DPa, DPb, and DPc. For decreasing values of  $\tau$  ([s]), the table contains the verification time ([min]), the success probability CI estimated through Uppaal, the mean success rate observed at runtime, the estimated maximum fatigue value for all humans, and the estimated minimum charge value for the robot. Fatigue and charge level are estimated for the maximum value of  $\tau$  for each scenario. For each metric, configurations leading to the least accurate results are highlighted in grey.

SC.	DSL2SHA	$\tau$	Ver.Time [min]	Success Probability		HUM.	Max. Fatigue		ROB.	Min. Charge	
				DT	DEPL ( $\bar{p}$ )		DT	DEPL		DT	DEPL
DPa	235/863 (27.2%)	400	16.36	0.933 $\pm$ 0.05	1.00	P1 D1	0.2664 $\pm$ 0.014	0.2385	Tbot	82.4%	82.77%
		350	54.67	0.706 $\pm$ 0.05	0.60		0.0372 $\pm$ 0.004	0.0332			
		300	59.09	0.419 $\pm$ 0.05	0.40						
DPb	235/876 (26.8%)	520	22.69	0.909 $\pm$ 0.05	1.00	P1 D1	0.2469 $\pm$ 0.009	0.2191	Tbot	80.4%	79.99%
		450	64.29	0.597 $\pm$ 0.05	0.50		0.0248 $\pm$ 0.007	0.0235			
		400	26.88	0.227 $\pm$ 0.05	0.20						
DPc	283/1161 (15.7%)	1500	54.54	0.920 $\pm$ 0.05	1.00	P1 D1 P2 D2	0.2860 $\pm$ 0.063	0.3070	Tbot	64.3%	67.85%
		1400	123.52	0.792 $\pm$ 0.05	0.80		0.0064 $\pm$ 0.002	0.0067			
		1300	175.68	0.421 $\pm$ 0.05	0.40		0.6028 $\pm$ 0.044	0.6610			
							0.0218 $\pm$ 0.002	0.0261			

been investigated with scenario LB), and patients exhibit more critical fatigue profiles than doctors.

As per Fig. 2, the entry point to the design-time phase analysis is the specification of the scenarios through the DSL presented in Section 3. The complete DSL file for the three scenarios is reported in Appendix B. All scenarios are set in the same layout (shown in Fig. 12), thus, there is only one floor definition. Agents participating in the three scenarios are fixed; specifically, the DSL features one robot definition (identified as Tbot) and eight human definitions (P1 and D1 for DPa, P1 and D1 for DPb, and P1, D1, P2, and D2 for DPc). We recall that the maximum velocity and acceleration for the robot are directly derived from its type parameter, which, in this case, is turtlebot3\_wafflepi. Each scenario in Table 6 corresponds to a robotic mission, thus, there are three mission definition blocks defining the sequence of services that the robot must provide to complete the mission with success. Finally, queries are defined to compute the metrics required to carry out this design-time analysis, i.e., the probability of success for decreasing values of  $\tau$ , estimated fatigue for all human subjects, and residual battery charge. Parameter R (the bound on runs) is set to auto, to indicate that Uppaal should generate as many runs as necessary to compute estimations with the requested confidence level. As per Fig. 5, for each mission (thus, in our case, DPa, DPb, DPc), a JSON file is automatically generated and converted into a pair of Uppaal model/query files to perform verification.

We assess the “efficiency” of the DSL in terms of effort saved compared to manually drafting the SHA network modeling each scenario. To this end, we calculate the ratio (indicated as DSL2SHA in Table 7) between the size of a DSL instance and the size of the corresponding SHA network. We compute the size of a DSL model as the number of words needed to configure the scenario. Counting words rather than abstract elements captured by the DSL gives us a more accurate indication of the DSL’s verbosity: note that, since the declaration of each element in Section 3.2 requires at least one word, counting abstract elements would result in more favorable ratios. Given a SHA  $\mathcal{A}$ , we compute its size, indicated as  $|\mathcal{A}|$  according to Eq.13:

$$|\mathcal{A}| = |\mathcal{E}| + |\Gamma(W)| + |C_{!?}| + |\Xi(W)| + |L| + |\mathcal{D}| + |\mathcal{F}| + |W| \quad (13)$$

The size of a network of SHA equals the sum of the sizes of all the SHA that compose it. Table 7 reports the resulting ratios.

SMC results are reported in Table 7 and discussed in the following. For this validation phase, the duration of verification experiments performed through Uppaal ranges from 16.36min to 175.68min in the worst case (i.e., scenario DPc, which has the most complex robotic mission and highest  $\tau$  values). Unlike the previous phase, the goal in this case is to test framework’s efficacy when developing realistic scenarios. Therefore, we do not collect a large batch of deployment traces to keep the duration of the deployment phase more practical (i.e., shorter than 1h). Given the smaller number of deployment traces that have been collected, the probability of success of the deployed system is not calculated through Algorithm 1, as it would yield

1202 scarcely significant CIs. In this case, we adopt *point estimator*  $\bar{p}$  given by the percentage of successful runs as  
 1203 specified by Eq.14, where  $\mathcal{DT}$  is the set of deployment traces and set  $\mathcal{SVD}_{dt}$  is calculated from a deployment  
 1204 trace  $dt \in \mathcal{DT}$  as in Algorithm 1, Line 3.

$$\bar{p} = \frac{|\{dt|dt \in \mathcal{DT} \wedge |\mathcal{SVD}_{dt}| = N_h \wedge \max(\mathcal{SVD}_{dt}) \leq \tau - T_{\text{int}}\}|}{|\mathcal{DT}|} \cdot 100 \quad (14)$$

1205 Metrics related to fatigue (for each human subject) and battery charge are computed as in the previous  
 1206 validation phase.

1207 Through the design-time analysis we estimate that the three scenarios require a  $\tau$  of approximately 7min,  
 1208 9min, and 25min, respectively, to end in success with probability greater than 90%. As previously mentioned,  
 1209 the robot’s charge is not critical for any of the scenarios: although DPc is the most demanding in terms of  
 1210 robot’s power, since the initial charge  $C_0$  is 99% the estimated residual charge at the end of the mission is  
 1211 greater than 60%. Doctors (i.e., agents D1 and D2) all adhere to the Elderly/Healthy fatigue profile, thus  
 1212 they do not constitute a criticality to the mission. As per Table 7, they reach an estimated maximum fatigue  
 1213 level between 2.18% and 3.7%, in particular D1 in DPc (who participates in the Recipient pattern) reaches  
 1214 the lowest fatigue value (0.6%) as they only move haphazardly out of free will while waiting for the robot  
 1215 to deliver the resource (see Section 4.1.3). On the other hand, as expected, patients reach more critical  
 1216 values. We remark that, although they walk for longer, patient P1 in all three scenarios reaches fatigue levels  
 1217 compatible with those estimated for HL because they have time to rest while the robot is assisting the doctor.

1218 The design-time analysis highlights that the most concerning aspect among the three scenarios is the  
 1219 fatigue level reached by patient P2 in DPc, as they also adhere to a critical fatigue profile (Elderly/Sick) and  
 1220 have to cover a significant distance from R1b to OFF3. For all experiments, threshold  $F_{\text{high}}$  (see Table 3) is  
 1221 set to 0.6. Therefore, some traces of the formal model feature the orchestrator instructing P2 to stop and  
 1222 rest (when  $f \geq F_{\text{high}}$  holds) causing a delay in the mission. We remark that this safety measure embedded in  
 1223 the orchestrator is necessary to prevent the patient from reaching the maximum value of fatigue, but it is  
 1224 not sufficient to prevent them from reaching a significant (average) fatigue level (i.e., approximately 60%).

1225 Results observed by deploying the scenarios in the hybrid setting corroborate the outcomes predicted at  
 1226 design-time. As in the first validation phase, Table 7 highlights the results corresponding to larger estimation  
 1227 errors. Specifically, success probability ranges estimated at design time and reported in column DT (we recall  
 1228 that rates in column DEPL are point estimators and, thus, not reported as ranges) are the least accurate  
 1229 when the average success rate is closer to 50% or 60% (DPa with  $\tau = 350s$  and DPb with  $\tau = 450s$ , also  
 1230 highlighted in grey). On the other hand, estimations of the fatigue level have design time-to-deployment  
 1231 differences range from approximately 5% in the best case (D1 in DPc) to 16% in the worst case (D2 in DPc,  
 1232 also highlighted in grey). Nevertheless, we recall that, although errors are larger than those obtained with the  
 1233 first three scenarios, these are not an indication of inaccuracies within the formal model as only 5 deployment  
 1234 traces are performed for DPa, DPb, and DPc (compared to more than 100 for the previous validation phase).

1235 Given the results of the first design-time analysis round and the data collected during deployment, the  
 1236 designer in charge of developing and maintaining these scenarios may choose to apply reconfiguration measures  
 1237 and refine the three robotic missions as described in Section 5. The reconfiguration measures applied to the  
 1238 three scenarios (hereinafter referred to as R-DPa, R-DPb, and R-DPc) and the updated sequences of services  
 1239 are described in Table 8. Since the robot’s battery was not a critical element in the first round of analysis,  
 1240 replacing the robot with a different one or recharging it would not impact the updated results. For scenarios  
 1241 DPa and DPb, the sequence of services (i.e., the robot’s mission) is modified to reduce the time required to  
 1242 complete the mission or, in other words, to obtain a high probability of success for smaller values of  $\tau$ . In  
 1243 these two cases, the patient is led straight to the examination room rather than to the waiting room first.<sup>10</sup>  
 1244 As for the third scenario, the goal is to lighten the strain on the patient in the most delicate condition.  
 1245 Therefore, the robot serves P2 first and leads them to the doctor’s office last to allow them a longer recovery  
 1246 time while in the waiting room. Table 9 reports the DSL2SHA ratio for the reconfigured scenarios. Note that

<sup>10</sup>Note that, in a real healthcare facility, this may not be feasible in all cases: the examination room must either be empty when P1 is served or equipped to host more than one patient simultaneously.

Table 8: Reconfiguration measures applied to scenarios DPa, DPb, and DPc, and updated sequence of services.

SCENARIO	RECONFIGURATION MEASURES	MISSION
R-DPa	The robot (Tbot) <i>leads</i> P1 directly to R2, then it serves D1 by <i>following</i> them to CUP1 and back to R2.	P1 Follower ENTR → R2, D1 Leader R2 → CUP1, D1 Leader CUP1 → R2
R-DPb	The robot (Tbot) serves D1 first by <i>fetching</i> the resource from CUP2 and <i>follows</i> them to R2, then it serves P1 and <i>leads</i> them to R2.	D1 Recipient OFF2 ↔ CUP2, D1 Leader OFF2 → R2, P1 Follower ENTR → R2
R-DPc	The robot (Tbot) <i>leads</i> P2 to R1b first and then provides the same sequence of services as scenario DPc.	P2 Follower ENTR → R1b, P1 Follower ENTR → R1a, D1 Recipient OFF1 ↔ CUP1, D2 Leader OFF3 → CUP2, D2 Leader CUP2 → OFF3, P1 Follower R1a → OFF1, P2 Follower R1b → OFF3

Table 9: Results of DSL2SHA calculation and of the design-time analysis of scenarios R-DPa, R-DPb, and R-DPc. For decreasing values of  $\tau$  ([s]), the table contains the verification time ([min]), the success probability CI estimated through Uppaal, the estimated maximum fatigue value for all humans, and the estimated minimum charge value for the robot. Fatigue and charge level are only estimated for the maximum value of  $\tau$  for each scenario.

SC. ( $M$ )	DSL2SHA	$\tau$	Ver.Time [min]	Success Probability ( $\mathbb{P}_M(\diamond_{\leq \tau} \text{scs})$ )	HUM.	Max. Fatigue ( $E_{M, \max(\tau)}[\max(F_i)]$ )	ROB.	Min. Charge ( $E_{M, \max(\tau)}[\min(C)]$ )
R-DPa	227/768 (29.5%)	300	6.86	$0.950 \pm 0.05$	P1 D1	$0.1042 \pm 0.014$ $0.0367 \pm 0.004$	Tbot	86.51%
		250	32.56	$0.498 \pm 0.05$				
		200	30.10	$0.258 \pm 0.05$				
R-DPb	227/781 (29.0%)	350	10.54	$0.950 \pm 0.05$	P1 D1	$0.1387 \pm 0.008$ $0.0235 \pm 0.001$	Tbot	85.45%
		320	55.46	$0.741 \pm 0.05$				
		300	40.69	$0.206 \pm 0.05$				
R-DPc	283/1161 (15.7%)	1500	20.60	$0.950 \pm 0.05$	P1 D1 P2 D2	$0.3034 \pm 0.061$ $0.0032 \pm 0.001$ $0.3761 \pm 0.029$ $0.0245 \pm 0.016$	Tbot	64.28%
		1400	121.59	$0.854 \pm 0.05$				
		1300	149.31	$0.556 \pm 0.05$				

1247 for R-DPc the ratio is unvaried since only the order in which humans are served is changed. On the other  
1248 hand, for R-DPa and R-DPb removing one human declaration (13 words) reduces the SHA network size by  
1249 11% and 10.8%, respectively.

1250 The results of the second round of design-time analysis are reported in Table 9. Quality metrics report a  
1251 slight improvement compared to the first round of analysis since the robotic mission is shorter for R-DPa  
1252 and R-DPb. In this case, verification time ranges from 10.54min to approximately 149.31min in the worst  
1253 case, as scenario R-DPc is substantially unvaried in terms of performance. Estimations inform us that for  
1254 R-DPa and R-DPb the updated mission can be completed successfully in less time as we obtain a success  
1255 probability  $> 90\%$  for  $\tau$  equal to 300s and 350s, respectively (compared to 400s and 520s for the initial  
1256 configuration). Since the patient only walks from the entrance to R2, their estimated maximum level of  
1257 fatigue is also approximately 60% (in R-DPa) and 43% (in R-DPb) lower than fatigue estimations obtained  
1258 with DPa and DPb, respectively, whereas the value remains essentially unchanged for D1 as they perform the  
1259 same actions as in the original scenario. For R-DPc, we observe that allowing P2 more time to rest in the  
1260 waiting room reduces their maximum fatigue level by 37%. Furthermore, as they do not reach the critical  
1261 threshold  $C_{\text{high}} = 60\%$  anymore, the orchestrator does not instruct them to stop mid-service, leading to a

1262 slight reduction in the duration of the mission.

### 1263 6.3. Discussion

1264 We can summarize how we have addressed the validation goals as follows:

1265 **G1.** We have performed more than 300 runs of three experimental scenarios in a digital-twin environment  
1266 involving simulated humans and a real robotic device communicating via ROS. Collected deployment  
1267 traces have been exploited to assess the accuracy of the formal model and SMC results.

1268 **G2.** We have assessed the coverage of our development framework with respect to existing real-world  
1269 scenarios in the service robotics domain. We have then collected the most recurring tasks within  
1270 the collected set of real applications into three scenarios to be analysed and developed through our  
1271 framework. In this regard:

- 1272 (a) We have assessed the efficiency of the presented DSL by calculating the number of words necessary  
1273 to configure the whole SHA network (i.e., the DSL2SHA metric).
- 1274 (b) We have analysed the three scenarios at design-time and the results of such analysis are reflected  
1275 by the deployment traces.
- 1276 (c) We have reconfigured the three scenarios in light of the collected deployment traces and iterated  
1277 the design-time analysis.

1278 Concerning the analysis of the formal model accuracy (goal **G1**), as discussed in Section 6.1, we obtain  
1279 relative estimation errors for the probability of success and charge level smaller than 10% also in boundary  
1280 conditions, e.g., involving subjects with a critical fatigue profile or a robot close to full discharge. Success  
1281 probability and minimum battery charge ranges provide empirical evidence of the reliability of the SHA  
1282 modeling the robotic system. Since we have only performed experiments with virtual human agents whose  
1283 model derives from literature analysis, the validation of the formal model of human behavior needs further  
1284 investigation. As future work, the validation process is to be completed by performing experiments with real  
1285 human subjects to assess the accuracy of SHA modeling human behavior.

1286 Coverage analysis (enabling the pursuit of goal **G2**) yields that more than 80% of the collected real-world  
1287 scenarios within the scope of this work can be designed and deployed through the presented framework.  
1288 The analysis carried out on scenarios DP<sub>a</sub>, DP<sub>b</sub>, and DP<sub>c</sub> shows how the framework supports practitioners  
1289 throughout the entire development process by automating the generation of the formal model and the  
1290 deployment of the resulting application.

1291 The analysis of the DSL2SHA ratio (smaller than 30% in all cases) shows that the DSL requires less  
1292 effort than manually drafting the formal model (goal **G2a**). DSL2SHA values show that the DSL grows more  
1293 efficient than the manual creation of the formal model as the size of the SHA network in question increases.  
1294 This is due to the fact that the portion of DSL configuring the geometrical layout (which is the same for  
1295 all scenarios in Section 6.2, regardless of the complexity of the mission) is the most verbose element. This  
1296 issue is to be addressed in the future by automatically acquiring the information regarding the layout from  
1297 planimetries to significantly boost the efficiency of the DSL.

1298 Indicators estimated through the design-time analysis phase of the three scenarios are corroborated by  
1299 the observations collected during deployment (goal **G2b**). With a small number of deployment traces (i.e.,  
1300 5), relative estimation errors do not exceed 16%. As for goal **G2c**, reconfiguration measures applied to  
1301 scenarios DP<sub>a</sub> and DP<sub>b</sub> (through minor modifications to the DSL specification) improve the estimated success  
1302 probabilities with a time bound smaller by 25% (300s compared to 400s) and 33% (350s compared to 570s),  
1303 respectively. As previously discussed, reconfiguring DP<sub>c</sub> reduces the physical effort imposed on subject P2.

## 1304 7. Related Work

1305 Introducing formal analysis into the robot software development process is a long-standing issue in the  
1306 research community. In a survey from 2019, Luckcuck et al. [53] examine more than 60 papers focused on

1307 specification and verification of autonomous robotic systems, emphasizing both the community’s interest in  
1308 the topic and the challenges to face as we move forward. In the following, we report on works existing in the  
1309 literature proposing:

- 1310 1. formal modeling and verification techniques for the analysis of robotic applications, especially dealing  
1311 with human-robot interaction;
- 1312 2. user-friendly DSLs for the specification of robotic missions.

### 1313 7.1. Formal Analysis for Robotic Applications Development

1314 Developing software for the robotic domain is an elaborate process given the complexity and unstructured  
1315 nature of the system itself [7]. Therefore, it usually requires a combination of different software development  
1316 techniques to achieve a satisfactory result. Several works focus on tasks such as testing and simulation [54] or  
1317 implementation [55], which are substantial to the development process but out of the scope of this review. In  
1318 the following, we focus on the early design phase and report on works exploiting formal methods to this end.

1319 Existing works can be classified based on the *formalism* used to model the environment and the agents’  
1320 behavior and the *verification* technique applied to check properties.

#### 1321 7.1.1. Temporal Logic-based Robotic Applications Modeling

1322 As for the first criterion, temporal logic notations are often adopted to model the robotic task. Gainer et  
1323 al. [56] present the CRutoN tool to analyse a personal robot’s behavior in a domestic setting. The work  
1324 models the robot’s behavior as a set of logic constraints, which are automatically parsed and converted into a  
1325 NuSMV model [57]. The generated model is put through model checking to verify relevant properties about  
1326 the system, e.g., that the robot never fails to alert the user about an event that requires their attention.  
1327 Webster et al. [58] had previously exploited the *BrahmsToPromela* tool [59] for the same case study. The  
1328 human users and the robot are modeled as *agents* using Brahms. Brahms models are then automatically  
1329 translated into Promela and verified through the SPIN model-checker. Both works treat human behavior  
1330 as a black-box whose actions are selected non-deterministically out of a pre-determined set. The work by  
1331 Vicentini et al. [60] introduces an innovative risk assessment procedure for collaborative industrial tasks  
1332 based on the TRIO temporal logic language [61]. Similarly to previous examples, the authors model the  
1333 agents and the task through a set of logic formulae to find safety hazards and assess their severity. As  
1334 previously mentioned, human-robot interaction introduces uncertainties into the model, thus, the work has  
1335 been subsequently extended to include manifestations of erroneous human behavior [62].

#### 1336 7.1.2. State-based Robotic Applications Modeling

1337 State-based formalisms are also a popular choice to model the behavior of robotic systems. Most works  
1338 pair the state-based model of the system with a set of logic properties to perform verification. Ding et al.  
1339 [63] exploit Finite State Machines to model collaborative industrial tasks, later extended to cover multi-robot  
1340 multi-human tasks [64], where unexpected events due to the presence of humans are modeled as exceptions  
1341 and paired with a recovery strategy. Porfirio et al. [65] explore how formal verification can be used to  
1342 ensure that robots adhere to *social* norms while interacting with humans. Norms, expressed as LTL formulae,  
1343 constitute the properties to be verified, whereas interaction sequences are modeled as a composition of  
1344 Labelled Transition Systems (LTSs). The work by Adam et al. [66] also targets the social robotics field, as  
1345 the authors propose the CAIO framework. The authors exploit the Belief Desire Intention (BDI) architecture  
1346 and models of human cognition to develop a perception and deliberation process that drives the robot towards  
1347 making decisions in a human-like fashion and making human-robot interaction feel more natural. Araiza-Illan  
1348 et al. [67] exploit the AgentSpeak language [68] to implement BDI agents and automatically generate test  
1349 cases for interactive robotic applications. The framework is tested on a cooperative table assembly case  
1350 study, where the robot’s BDI agent infers the human’s state based on three sensors and reacts accordingly as  
1351 encoded by the AgentSpeak model. Quottrup et al. [69] model multi-robot systems as a network of Timed  
1352 Automata and verify whether collisions potentially occur or some robots are not able to complete their goal,  
1353 which are all expressed as CTL properties and verified through Uppaal. Zhou et al. [70] propose a similar



1354 approach based on Timed Automata and MITL properties focused on motion planning to synthesize optimal  
1355 trajectories based on verification results. Some works have also exploited Hybrid Automata to incorporate  
1356 physical laws into the verification process. Molnar et al. [71] introduce the concept of Model Composition  
1357 Agents (MCA), which encapsulate a Hybrid Automaton modeling either an agent or the environment and its  
1358 interaction with other automata in the system. The resulting network of MCA is abstracted as an LTS and  
1359 model checked to diagnose faults in the original system.

### 1360 7.1.3. Formalizations of Human Behavior

1361 As human-robot interaction becomes a key element in modern robotic systems, particular attention has  
1362 been given to how the unpredictability due to the presence of humans can be formally modeled. In this  
1363 aspect, two main research directions emerge from the literature: game-based approaches and probabilistic  
1364 models. The possibility to model the interaction between a robotic agent and the environment as a *game* to  
1365 synthesize a robot controller strategy (if it exists) is investigated in [72]. Kress et al. emphasize the challenge  
1366 to find a proper abstraction of the environment model that allows for significant verification results without  
1367 leading to state space explosion. Chen et al. [73] apply the approach based on Timed Game Automata  
1368 (TGA) and LTL to surveillance, monitoring, and delivery tasks in partially unknown environments. The work  
1369 by Bersani et al. [74] addresses applications involving robots and humans working in a shared environment,  
1370 modeled as TGA networks. Humans are modeled as *uncontrollable* agents, to capture the uncertainty of  
1371 their behavior. A robot controller that also accounts for unpredictable human moves is then automatically  
1372 synthesized through the Uppaal-TIGA tool.

1373 On the other hand, probabilistic models of human behavior and decision making (e.g., the *Boltzmann*  
1374 *policy* [75]) are well-established in the literature and have been successfully applied to the robotic domain.  
1375 Mason et al. [76] exploit Markov Decision Processes (MDPs) to model an assistive-living scenario and verify  
1376 probabilistic properties (expressed in PCTL logic) through the PRISM model checker [77]. The work by  
1377 Junges et al. [78] combines the two approaches since it models the robot as a stochastic *controllable* agent  
1378 and the human as stochastic and *uncontrollable*, which, when combined, produce a stochastic two-player  
1379 game. In this case, optimal robot policies are also synthesized through PRISM-Games [79]. Vibekananda et  
1380 al. [80] exploit Probabilistic State Machines to perform human pose estimation and predict their intention  
1381 while interacting with a robot. Galin et al. [81] build upon a previous study on how Cellular Automata with  
1382 probabilistic transitions can be used to model human motion in partially unknown environments [82]. The  
1383 authors exploit these theoretical results to develop the model of a shared workspace where human and robot  
1384 work simultaneously to compute the area where their trajectories are more likely to overlap.

### 1385 7.1.4. Verification Techniques and Tools

1386 Since state-based formalisms and temporal logics are the most popular choices when it comes to modeling  
1387 the robotic system, it follows that model-checking is the natural choice in terms of verification technique [53],  
1388 given the availability of powerful model checkers such as Uppaal [27] and SPIN [83]. Models based on MDPs,  
1389 such as the one developed by Ye et al. [84], can be verified through Probabilistic Model Checking, which  
1390 is most often performed through PRISM [77]. Statistical Model Checking (SMC), which is the verification  
1391 technique used in our framework, has also gained momentum over the last few years. The most common  
1392 motivation pertains to the reduced verification times, which lead to more practical approaches. Paigwar et al.  
1393 [85] exploit SMC to estimate the probability of collisions in automated driving systems. Foughali et al. [86]  
1394 apply SMC to formally verify real-time properties, like schedulability and readiness, of robotic software. Herd  
1395 et al. [87] focus on multi-agent systems, and on swarm robotics in particular: in this case, SMC dampens  
1396 issues related to the size of the problem, which cannot be handled by traditional model checking techniques.

## 1397 7.2. Specification Languages for the Robotic Domain

1398 In 2014, Nordmann et al. [88] surveyed 137 papers presenting robotic DSLs. At the time of writing,  
1399 Scopus indexes more than 90 papers published since 2014 with keywords *robot\** and *domain-specific language*.  
1400 These numbers show that DSL development is a cornerstone of the robotic software engineering process since  
1401 it automates the generation of code or complex models and makes development frameworks accessible to a

1402 wider audience. Referring to the classification in [88], in the following we report on the subset of works on  
1403 this topic dealing with the *scenario building* phase, i.e., DSLs to specify high-level environment features and  
1404 the robot’s task, as these are the closest to our work.

### 1405 7.2.1. DSLs for Scenario Building.

1406 Noreils and Chatila [89] present a high-level notation to specify reactive robotic mission plans. The  
1407 language envisages the specification of modules, which are further structured into three architectural layers:  
1408 the *functional* layer to specify the lower-level robot’s capabilities, the *planning* layer to specify task sequences,  
1409 and the *control* layer that translates plans into requests to the functional modules. Knoop et al. [90] present  
1410 an approach to automatically generate robotic tasks starting from representations of tasks in the human  
1411 operational space, adhering to the Programming by Demonstration paradigm. Finucane et al. [91] present  
1412 the LTLMoP framework to automatically synthesize and deploy robot controllers. The framework converts  
1413 Structured English specifications describing the robotic task into equivalent LTL formulae, which are then  
1414 synthesized into an automaton (the *discrete* controller). The work has been subsequently extended by Raman  
1415 et al. [92] with implicit memory strategies to model robotic tasks depending on events that occurred in the  
1416 past (e.g., “every time you sense *order*, visit the *kitchen*”). Kunze et al. [93] present SRDL, a framework  
1417 extending the KnowRob knowledge base [94] with notions about *robots*, hardware *components*, *actions*, and  
1418 *capabilities* (of performing a certain action). Miyazawa et al. [95] introduce RoboChart, a DSL to model  
1419 and verify real-time concurrent robotic tasks with budgets and deadlines (i.e., cost and time constraints).  
1420 RoboChart semantics, which is based on Timed Automata and Timed Communicating Sequential Processes  
1421 (CSP) [96], makes the notation amenable to formal verification, specifically model-checking. Ciccuzzi et al.  
1422 [97] propose a family of three languages to specify missions for multi-robot systems: the Monitoring Mission  
1423 Language to specify task sequences, the Robot Language to configure the individual robots, and the Behavior  
1424 Language to specify the atomic movements of robots.

### 1425 7.2.2. DSLs for Human-Robot Interaction

1426 The advent of human-robot interaction and collaborative robotics has also influenced DSL development.  
1427 Gavran et al. [98] introduce the TOOL DSL to specify collaborative assembly tasks in the manufacturing  
1428 sector through a textual notation that is accessible to non-experts. Detzner et al. [99] present LoTLan, a  
1429 DSL to describe material flow processes in warehouses. The work consists of a procedure to map human vocal  
1430 requests (e.g., “I need an item”) to common semantics, identifying *who* has to perform *which action*, and  
1431 finally LoTLan primitives, which are then converted into plans for AGVs. Forbig et al. [100] exploit their  
1432 language CoTaL [101] to model interactive tasks between a humanoid robot and a stroke patient performing  
1433 arm mobility recovery exercises. The resulting specification captures all phases needed for the exercise session,  
1434 how the humanoid robot can detect whether the patient has completed an exercise or not and how to react  
1435 accordingly.

### 1436 7.3. Discussion

1437 As this survey shows, numerous approaches exploit formal methods to analyze robotic applications.  
1438 Specifically, several attempts have been made at formalizing the aspects of human behavior that are  
1439 significant while interacting with a robot and should, thus, impact the results of the formal analysis. Most of  
1440 these works present approaches that are either deterministic, game-based, or probabilistic, such as the hereby  
1441 presented framework. Deterministic approaches—such as architectures based on BDI agents—potentially  
1442 result in less complex models and more favorable verification times. However, assuming complete rationality  
1443 and absence of fuzziness is reasonable for robotic agents (the orchestrator SHA indeed inherits most of its  
1444 substructures from the BDI architecture) or for human agents performing small repetitive tasks in controlled  
1445 environments [67, 66]. Human-robot interactions in the service sector feature virtually no constraint on  
1446 human behavior, thus deterministic models are overly restrictive. Furthermore, service robots applications  
1447 involve people from various age groups with different characteristics and performing a broad range of tasks.  
1448 Therefore, while estimating the outcome of a scenario, the exploration of the state space of all possible  
1449 behaviors should be guided by such features. Game-based approaches, although effective when exploited

1450 for controller synthesis [74], imply an exhaustive exploration of human actions (i.e., the *opponent's* move)  
1451 irrespective of their likelihood given the specific scenario configuration. For these reasons, probabilistic  
1452 approaches are particularly suited for the purpose of this framework. Specifically, to the best of the authors'  
1453 knowledge, this is the first attempt at combining probabilistic weights on human actions with a hybrid and  
1454 stochastic characterization of physiological processes. Due to its complexity, the resulting model is more  
1455 practically manageable through SMC rather than probabilistic model checking. Indeed, works exploiting  
1456 exhaustive techniques such as [102] focus on smaller setups targeting a specific task (e.g., the handover of an  
1457 item). Despite the loss in reliability introduced by SMC that only relies on a finite set of runs of the systems,  
1458 the proposed framework is applicable to a broad range of scenarios (as shown by the coverage analysis results)  
1459 while still providing results at design-time that accurately reflect runtime observations.

1460 As per Section 7.2, the literature is rich with DSLs for the robotic domain, but proposals targeting  
1461 interactive applications are lacking. Specifically, existing works target the manufacturing sector [98, 99] or  
1462 very specific tasks from the healthcare setting [100], whereas the service sector calls for more general-purpose  
1463 primitives to define how robots and humans interact. Other works propose a high-level specification of  
1464 mission patterns for multi-robot teams in environments (possibly) populated by humans [46, 12], but this has  
1465 not been attempted for applications where humans are actively involved as in the domain of this framework.

## 1466 8. Conclusion and Future Work

1467 We have presented a specification, modeling and analysis framework targeting interactive service robotic  
1468 applications. The framework has been extended with respect to previous publications with the introduction of  
1469 a custom DSL, refined formal models of the battery and the human behavior with a stochastic characterization  
1470 of human fatigue, and extensive experimental validation results including real-life experiments, coverage  
1471 analysis, and DSL evaluation.

1472 The framework is open to extensions. The quality metrics analyzed in the paper are not the only measures  
1473 of interest for a potential application designer using our modeling approach. With the current model, the  
1474 analysis can be easily enriched with measures such as the percentage of time agents spend in a certain  
1475 operational state, the frequency of a human agent reacting to a command issued by the orchestrator, the  
1476 number of times a certain action is taken, etc. For the sake of clarity, we kept the presentation limited to  
1477 the human fatigue and battery charge. In addition, a number of physiological or psychological indicators  
1478 can be considered to enrich the model of humans, provided that they can be represented by means of (OD)  
1479 equations or discrete/automata-based features. These indicators would allow the designer to model the  
1480 sensitivity of humans to phenomena that affect their responsiveness when environmental conditions are not  
1481 ideal. For instance, meaningful physiological values can be the heartbeat rate, the blood pressure, the breath  
1482 frequency of the patients in critical health conditions, and psychological indicators include stress, patience or  
1483 the level of engagement of the operators and doctors who take part in a scenario.

1484 The availability of tools such as the one presented in this work implies the need for a criterion to establish  
1485 which is the “right accuracy” to consider when judging the outcomes of the analysis. To the best of the  
1486 authors' knowledge, such a criterion does not exist and its definition would require specific work, possibly  
1487 conducted by healthcare specialists and medical engineers in real contexts. Nonetheless, the paper shows in  
1488 Section 6 an approach to assess the accuracy of our tool that is based on the comparison of quantitative  
1489 metrics obtained as a result of the formal analysis of the models and implemented real-world scenarios.

1490 We plan on improving the level of support provided to the designer by—at least partially—automating  
1491 the reconfiguration phase. The model of human behavior can be automatically refined based on observations  
1492 collected during deployment. The manual effort required on the designer's side can be further reduced by  
1493 automatically computing alternative mission plans leading to better key indicator values (success probability  
1494 and human subjects' physical strain) than those resulting from the initial plan. In addition, we plan on  
1495 assessing the DSL with the engagement of non-expert users of formal verification and of healthcare operators,  
1496 and to add syntactical structures that allow operators to integrate their own interaction patterns in the  
1497 framework without resorting to customized translations set out by formal method experts.

1498 **References**

1499 [1] C. B. Frey, M. A. Osborne, The future of employment: How susceptible are jobs to computerisation?, *Technological*  
1500 *forecasting and social change* 114 (2017) 254–280.

1501 [2] A. A. Morgan, J. Abdi, M. A. Syed, G. E. Kohen, P. Barlow, M. P. Vizcaychipi, Robots in healthcare: a scoping review,  
1502 *Current Robotics Reports* (2022) 1–10.

1503 [3] Robots and healthcare saving lives together (2022).  
1504 URL <https://www.automate.org/industry-insights/robots-and-healthcare-saving-lives-together>

1505 [4] Global healthcare assistive robot market (2022).  
1506 URL <https://www.globenewswire.com/news-release/2021/09/24/2303007/0/en/Global-Healthcare-Assistive-Robot-Market-Expected-19-CAGR-Will-Reach-USD-1-2-Billion-by-2026-Facts-Factors.html>

1507 [5] Robotics and the impact on nursing practice (2020).  
1508 URL [https://www.nursingworld.org/~494055/globalassets/innovation/robotics-and-the-impact-on-nursing-practice\\_print\\_12-2-2020-pdf-1.pdf](https://www.nursingworld.org/~494055/globalassets/innovation/robotics-and-the-impact-on-nursing-practice_print_12-2-2020-pdf-1.pdf)

1509 [6] Fraunhofer Institute for Manufacturing Engineering and Automation, EFFIROB: Economic feasibility studies on innovative  
1510 service robot applications (2010).  
1511 URL [https://www.ipa.fraunhofer.de/en/reference\\_projects/EFFIROB.html](https://www.ipa.fraunhofer.de/en/reference_projects/EFFIROB.html)

1512 [7] S. García, D. Strüber, D. Brugali, T. Berger, P. Pelliccione, Robotics software engineering: A perspective from the service  
1513 robotics domain, in: *ESEC/FSE, ACM, USA*, 2020, pp. 593–604.

1514 [8] K. Ehrlenspiel, A. Kiewert, U. Lindemann, M. S. Hundal, *Cost-efficient design*, Vol. 544, Springer, 2007.

1515 [9] D. Brugali, Software product line engineering for robotics, *Software Engineering for Robotics* (2021) 1–28.

1516 [10] P. Payne, M. Lopetegui, S. Yu, A review of clinical workflow studies and methods, in: *Cog. Inf.*, Springer, 2019, pp. 47–61.

1517 [11] ISO/PAS 21448:2019, Road vehicles — Safety of the intended functionality, ISO/PAS, 2019.

1518 [12] S. García, P. Pelliccione, C. Menghi, T. Berger, T. Bures, PROMISE: high-level mission specification for multiple robots,  
1519 in: *Intl. Conf. on Software Engineering, ACM, Seoul, South Korea*, 2020, pp. 5–8.

1520 [13] L. Lestingi, M. Askarpour, M. M. Bersani, M. Rossi, Formal verification of human-robot interaction in healthcare scenarios,  
1521 in: *SEFM, Springer*, 2020, pp. 303–324.

1522 [14] L. Lestingi, M. Askarpour, M. M. Bersani, M. Rossi, A model-driven approach for the formal analysis of human-robot  
1523 interaction scenarios, in: *IEEE SMC*, 2020, pp. 1907–1914.

1524 [15] L. Lestingi, M. Askarpour, M. M. Bersani, M. Rossi, A deployment framework for formally verified human-robot  
1525 interactions, *IEEE Access* 9 (2021) 136616–136635.

1526 [16] L. Lestingi, C. Sbrolli, P. Scarmozzino, G. Romeo, M. M. Bersani, M. Rossi, Formal modeling and verification of multi-robot  
1527 interactive scenarios in service settings, in: *Intl. Conf. on Formal Methods in Software Engineering*, 2022, pp. 80–90.

1528 [17] A. David, K. G. Larsen, A. Legay, M. Mikucionis, D. B. Poulsen, Uppaal SMC tutorial, *STTT* 17 (4) (2015) 397–415.

1529 [18] C. J. Clopper, E. S. Pearson, The use of confidence or fiducial limits illustrated in the case of the binomial, *Biometrika*  
1530 26 (4) (1934) 404–413.

1531 [19] R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, S. Yovine, The  
1532 algorithmic analysis of hybrid systems, *TCS* 138 (1) (1995) 3–34.

1533 [20] S. F. Arenis, M. Vujinovic, B. Westphal, On implementable timed automata, in: *Formal Techniques for Distributed*  
1534 *Objects, Components, and Systems*, Vol. 12136 of *Lecture Notes in Computer Science*, Springer, Valletta, Malta, 2020, pp.  
1535 78–95.

1536 [21] A. David, K. G. Larsen, A. Legay, M. Mikucionis, D. B. Poulsen, J. van Vliet, Z. Wang, Statistical model checking for  
1537 networks of priced timed automata, in: *Formal Modeling and Analysis of Timed Systems*, Vol. 6919 of *Lecture Notes in*  
1538 *Computer Science*, Springer, Aalborg, Denmark, 2011, pp. 80–96.

1539 [22] U. Grenander, Stochastic processes and statistical inference, *Arkiv för matematik* 1 (3) (1950) 195–277.

1540 [23] G. Agha, K. Palmkog, A survey of statistical model checking, *TOMACS* 28 (1) (2018) 1–39.

1541 [24] R. Alur, T. Feder, T. A. Henzinger, The benefits of relaxing punctuality, *Journal of the ACM (JACM)* 43 (1) (1996)  
1542 116–146.

1543 [25] D. Tardioli, R. Parasuraman, P. Ögren, Pound: A multi-master ros node for reducing delay and jitter in wireless multi-robot  
1544 networks, *Robotics and Autonomous Systems* 111 (2019) 73–87.

1545 [26] G. Behrmann, A. Cougnard, A. David, E. Fleury, K. G. Larsen, D. Lime, Uppaal-tiga: Time for playing games!, in:  
1546 *International Conference on Computer Aided Verification*, Springer, 2007, pp. 121–125.

1547 [27] K. G. Larsen, P. Pettersson, W. Yi, UPPAAL in a nutshell, *International Journal on Software Tools for Technology*  
1548 *Transfer* 1 (1-2) (1997) 134–152.

1549 [28] J. Z. Liu, R. W. Brown, G. H. Yue, A dynamical model of muscle activation, fatigue, and recovery, *Biophysical Journal*  
1550 82 (5) (2002) 2344–2359.

1551 [29] S. Konz, Work/rest: Part ii-the scientific basis (knowledge base) for the guide 1, *EGPS* 1 (401) (2000) 38.

1552 [30] Z. Givi, M. Y. Jaber, W. P. Neumann, Modelling worker reliability with learning and fatigue, *Applied Mathematical*  
1553 *Modelling* 39 (17) (2015) 5186–5199.

1554 [31] B. Liu, L. Ma, C. Chen, Z. Zhang, Experimental validation of a subject-specific maximum endurance time model,  
1555 *Ergonomics* 61 (6) (2018) 806–817.

1556 [32] M. Roberto, D. Farina, M. Gazzoni, M. Schieronni, Effect of age on muscle functions investigated with surface electromyog-  
1557 raphy, *Muscle & nerve* 25 (1) (2002) 65–76.

1558 [33] M. Hadley, A deterministic model of the free will phenomenon, *Journal of Consciousness Exploration & Research* 9 (1)  
1559 (2018).

- 1562 [34] C. Calude, F. Kroon, N. Poznanovic, Free will is compatible with randomness, *Philosophical Inquiries* 4 (2) (2016) 37–52.
- 1563 [35] M. Lutz, D. Stampfer, A. Lotz, C. Schlegel, Service robot control architectures for flexible and robust real-world task
- 1564 execution: Best practices and patterns, in: *INFORMATIK 2014*, Vol. P-232 of LNI, GI, Stuttgart, Germany, 2014, pp.
- 1565 1295–1306.
- 1566 [36] O. Tremblay, L.-A. Dessaint, A.-I. Dekkiche, A generic battery model for the dynamic simulation of hybrid electric
- 1567 vehicles, in: *Vehicle Power and Propulsion Conference*, IEEE, 2007, pp. 284–289.
- 1568 [37] H. Chuangfeng, L. Pingan, J. Xueyan, Measurement and analysis for lithium battery of high-rate discharge performance,
- 1569 *Procedia Engineering* 15 (2011) 2619–2623.
- 1570 [38] R. Merletti, L. L. Conte, C. Orizio, Indices of muscle fatigue, *J. of Electromyography and Kinesiology* 1 (1) (1991) 20–33.
- 1571 [39] L. Lindstrom, R. Kadefors, I. Petersen, An electromyographic index for localized muscle fatigue, *Journal of Applied*
- 1572 *Physiology* 43 (4) (1977) 750–754.
- 1573 [40] A. Rasheed, O. San, T. Kvamsdal, Digital twin: Values, challenges and enablers from a modeling perspective, *IEEE*
- 1574 *Access* 8 (2020) 21980–22012.
- 1575 [41] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng, ROS: an open-source robot operating
- 1576 system, in: *ICRA Workshop on Open Source Software*, Vol. 3, IEEE, Kobe, Japan, 2009, p. 5.
- 1577 [42] T. Nguyen, M. Reynolds, R. Kandaswamy, et al., Emerging technologies and trends impact radar: 2021, Gartner Research
- 1578 Notes. Available at (2021).
- 1579 [43] H. G. Kang, J. B. Dingwell, Differential changes with age in multiscale entropy of electromyography signals from leg
- 1580 muscles during treadmill walking, *PLoS one* 11 (8) (2016) e0162034.
- 1581 [44] E. T. Jaynes, O. Kempthorne, Confidence intervals vs bayesian intervals, in: *Foundations of probability theory, statistical*
- 1582 *inference, and statistical theories of science*, Springer, 1976, pp. 175–257.
- 1583 [45] F. Scholz, Confidence bounds and intervals for parameters relating to the binomial, negative binomial, poisson and
- 1584 hypergeometric distributions with applications to rare events, 2008 (2008).
- 1585 [46] C. Menghi, C. Tsigkanos, P. Pelliccione, C. Ghezzi, T. Berger, Specification patterns for robotic missions, *IEEE Trans.*
- 1586 *Software Eng.* 47 (10) (2021) 2208–2224.
- 1587 [47] M. Askarpour, C. Tsigkanos, C. Menghi, R. Calinescu, P. Pelliccione, S. García, R. Caldas, T. J. von Oertzen, M. Wimmer,
- 1588 L. Berardinelli, et al., RoboMAX: Robotic Mission Adaptation eXemplars, in: *2021 International Symposium on Software*
- 1589 *Engineering for Adaptive and Self-Managing Systems (SEAMS)*, IEEE, 2021, pp. 245–251.
- 1590 [48] K. Baraka, M. M. Veloso, Mobile service robot state revealing through expressive lights: formalism, design, and evaluation,
- 1591 *International Journal of Social Robotics* 10 (1) (2018) 65–92.
- 1592 [49] Case Studies - Service Robots, <https://ifr.org/case-studies/service-robots-case-studies> (2018).
- 1593 [50] M. Bajones, D. Fischinger, A. Weiss, P. D. L. Puente, D. Wolf, M. Vincze, T. Körtner, M. Weninger, K. Papoutsakis,
- 1594 D. Michel, et al., Results of field trials with a mobile service robot for older adults in 16 private households, *ACM*
- 1595 *Transactions on Human-Robot Interaction (THRI)* 9 (2) (2019) 1–27.
- 1596 [51] Care-O-bot 3 Application Scenarios, <https://www.care-o-bot.de/en/care-o-bot-3/application.html> (2012).
- 1597 [52] ISO 13482, Robots and robotic devices - Safety requirements for personal care robots, ISO, 2014.
- 1598 [53] M. Luckcuck, M. Farrell, L. A. Dennis, C. Dixon, M. Fisher, Formal specification and verification of autonomous robotic
- 1599 systems: A survey, *ACM Comput. Surv.* 52 (5) (2019) 100:1–100:41.
- 1600 [54] S. Abbaspour Asadollah, R. Inam, H. Hansson, A survey on testing for cyber physical system, in: *International Conference*
- 1601 *on Testing Software and Systems*, Springer, Sharjah and Dubai, UAE, 2015, pp. 194–207.
- 1602 [55] G. Ajaykumar, M. Steele, C.-M. Huang, A survey on end-user robot programming, *ACM Computing Surveys* 54 (8) (2021)
- 1603 1–36.
- 1604 [56] P. Gainer, C. Dixon, K. Dautenhahn, M. Fisher, U. Hustadt, J. Saunders, M. Webster, Cruton: Automatic verification of
- 1605 a robotic assistant's behaviours, in: *Critical Systems: Formal Methods and Automated Verification*, Springer, Turin, Italy,
- 1606 2017, pp. 119–133.
- 1607 [57] A. Cimatti, E. Clarke, E. Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani, A. Tacchella, Nusmv 2:
- 1608 An opensource tool for symbolic model checking, in: *International conference on computer aided verification*, Springer,
- 1609 Springer, Copenhagen, Denmark, 2002, pp. 359–364.
- 1610 [58] M. Webster, C. Dixon, M. Fisher, M. Salem, J. Saunders, K. L. Koay, K. Dautenhahn, Formal verification of an autonomous
- 1611 personal robotic assistant, in: *AAAI Spring Symposia*, AAAI Press, Palo Alto, California, USA, 2014, pp. 1–6.
- 1612 [59] R. Stocker, L. A. Dennis, C. Dixon, M. Fisher, Verifying brahms human-robot teamwork models, in: *Logics in Artificial*
- 1613 *Intelligence*, Vol. 7519 of *Lecture Notes in Computer Science*, Springer, Toulouse, France, 2012, pp. 385–397.
- 1614 [60] F. Vicentini, M. Askarpour, M. G. Rossi, D. Mandrioli, Safety assessment of collaborative robotics through automated
- 1615 formal verification, *IEEE Transactions on Robotics* 36 (1) (2019) 42–61.
- 1616 [61] C. A. Furia, D. Mandrioli, A. Morzenti, M. Rossi, Modeling Time in Computing, *Monographs in Theoretical Computer*
- 1617 *Science. An EATCS Series*, Springer, Berlin, Heidelberg, 2012.
- 1618 [62] M. Askarpour, D. Mandrioli, M. Rossi, F. Vicentini, Formal model of human erroneous behavior for safety analysis in
- 1619 collaborative robotics, *Robotics and Computer-Integrated Manufacturing* 57 (2019) 465–476.
- 1620 [63] H. Ding, J. Heyn, B. Matthias, H. Staab, Structured collaborative behavior of industrial robots in mixed human-robot
- 1621 environments, in: *Intl. Conf. on Automation Science and Engineering*, IEEE, Madison, WI, USA, 2013, pp. 1101–1106.
- 1622 [64] H. Ding, M. Schipper, B. Matthias, Collaborative behavior design of industrial robots for multiple human-robot collabora-
- 1623 tion, in: *IEEE International Symposium on Robotics*, IEEE, Seoul, Korea (South), 2013, pp. 1–6.
- 1624 [65] D. Porfirio, A. Sauppé, A. Albarghouthi, B. Mutlu, Authoring and verifying human-robot interactions, in: *ACM Symposium*
- 1625 *on User Interface Software and Technology*, ACM, Berlin, Germany, 2018, pp. 75–86.
- 1626 [66] C. Adam, W. Johal, D. Pellier, H. Fiorino, S. Pesty, Social human-robot interaction: A new cognitive and affective

- interaction-oriented architecture, in: *Intl. Conf. on Social Robotics*, Vol. 9979 of *Lecture Notes in Computer Science*, Springer, Kansas City, MO, USA, 2016, pp. 253–263.
- [67] D. Araiza-Illan, A. G. Pipe, K. Eder, Intelligent agent-based stimulation for testing robotic software in human-robot interactions, in: *Workshop on Model-Driven Robot Software Engineering*, ACM, Leipzig, Germany, 2016, pp. 9–16.
- [68] A. S. Rao, AgentSpeak(L): BDI agents speak out in a logical computable language, in: *Workshop on Modelling Autonomous Agents in a Multi-Agent World*, Vol. 1038 of *Lecture Notes in Computer Science*, Springer, Eindhoven, The Netherlands, 1996, pp. 42–55.
- [69] M. M. Quottrup, T. Bak, R. Izadi-Zamanabadi, Multi-robot planning: a timed automata approach, in: *IEEE International Conference on Robotics and Automation*, IEEE, New Orleans, LA, USA, 2004, pp. 4417–4422.
- [70] Y. Zhou, D. Maity, J. S. Baras, Timed automata approach for motion planning using metric interval temporal logic, in: *European Control Conference*, IEEE, Aalborg, Denmark, 2016, pp. 690–695.
- [71] L. Molnar, S. M. Veres, Hybrid automata discretising agents for formal modelling of robots, *IFAC Proceedings Volumes* 44 (1) (2011) 49–54.
- [72] H. Kress-Gazit, T. Wongpiromsarn, U. Topcu, Correct, reactive, high-level robot control, *IEEE Robotics & Automation Magazine* 18 (3) (2011) 65–74.
- [73] Y. Chen, J. Tumova, C. Belta, LTL robot motion control based on automata learning of environmental dynamics, in: *IEEE International Conference on Robotics and Automation*, IEEE, St. Paul, Minnesota, USA, 2012, pp. 5177–5182.
- [74] M. M. Bersani, M. Soldo, C. Menghi, P. Pelliccione, M. Rossi, PuRSUE—from specification of robotic environments to synthesis of controllers, *Formal Aspects of Computing* 32 (2) (2020) 187–227.
- [75] C. L. Baker, J. Tenenbaum, R. R. Saxe, Goal inference as inverse planning, in: *Proceedings of the Annual Meeting of the Cognitive Science Society*, 29 (29), 2007.
- [76] G. Mason, R. Calinescu, D. Kudenko, A. Banks, Assurance in reinforcement learning using quantitative verification, *Advances in hybridization of intelligent methods* 85 (2017) 71.
- [77] M. Z. Kwiatkowska, G. Norman, D. Parker, PRISM 4.0: Verification of probabilistic real-time systems, in: *Computer Aided Verification*, Vol. 6806 of *Lecture Notes in Computer Science*, Springer, Snowbird, UT, USA, 2011, pp. 585–591.
- [78] S. Junges, N. Jansen, J.-P. Katoen, U. Topcu, Probabilistic model checking for complex cognitive tasks—a case study in human-robot interaction, *arXiv preprint arXiv:1610.09409* (2016).
- [79] T. Chen, V. Forejt, M. Kwiatkowska, D. Parker, A. Simaitis, PRISM-games: A model checker for stochastic multi-player games, in: *Intl. Conf. on TOOLS and Algorithms for the Construction and Analysis of Systems*, Springer, Rome, Italy, 2013, pp. 185–191.
- [80] V. Dutta, T. Zielinska, Predicting the intention of human activities for real-time human-robot interaction (HRI), in: *Intl. Conf. on Social Robotics*, Vol. 9979 of *Lecture Notes in Computer Science*, Springer, Kansas City, MO, USA, 2016, pp. 723–734.
- [81] R. R. Galin, R. V. Meshcheryakov, M. V. Mamchenko, Analysis of intersection of working areas within the human-robot interaction in a shared workspace, in: *Proceedings of the Computational Methods in Systems and Software*, Springer, Czech Republic, 2021, pp. 749–759.
- [82] R. Breukelaar, T. Bäck, Using a genetic algorithm to evolve behavior in multi dimensional cellular automata: emergence of behavior, in: *Genetic and Evolutionary Computation Conference*, ACM, Washington DC, USA, 2005, pp. 107–114.
- [83] G. J. Holzmann, The model checker SPIN, *IEEE Transactions on Software Engineering* 23 (5) (1997) 279–295.
- [84] K. Ye, A. Cavalcanti, S. Foster, A. Miyazawa, J. Woodcock, Probabilistic modelling and verification using RoboChart and PRISM, *Software and Systems Modeling* (2021) 1–50.
- [85] A. Paigwar, E. Baranov, A. Renzaglia, C. Laugier, A. Legay, Probabilistic collision risk estimation for autonomous driving: Validation via statistical model checking, in: *IEEE Intelligent Vehicles Symposium*, IEEE, Las Vegas, NV, USA, 2020, pp. 737–743.
- [86] M. Foughali, F. Ingrand, C. Seceleanu, Statistical model checking of complex robotic systems, in: *International Symposium on Model Checking Software*, Vol. 11636 of *Lecture Notes in Computer Science*, Springer, Beijing, China, 2019, pp. 114–134.
- [87] B. Herd, S. Miles, P. McBurney, M. Luck, Quantitative analysis of multiagent systems through statistical model checking, in: *Intl. Workshop on Engineering Multi-Agent Systems*, Springer, 2015, pp. 109–130.
- [88] A. Nordmann, N. Hochgeschwender, S. Wrede, A survey on domain-specific languages in robotics, in: *Simulation, Modeling, and Programming for Autonomous Robots*, Vol. 8810 of *Lecture Notes in Computer Science*, Springer, Bergamo, Italy, 2014, pp. 195–206.
- [89] F. R. Noreils, R. Chatila, Plan execution monitoring and control architecture for mobile robots, *IEEE Trans. Robotics Autom.* 11 (2) (1995) 255–266.
- [90] S. Knoop, M. Pardowitz, R. Dillmann, Automatic robot programming from learned abstract task knowledge, in: *Intl. Conf. on Intelligent Robots and Systems*, IEEE, California, USA, 2007, pp. 1651–1657.
- [91] C. Finucane, G. Jing, H. Kress-Gazit, LTLMoP: Experimenting with language, temporal logic and robot control, in: *Intl. Conf. on Intelligent Robots and Systems*, IEEE, Taipei, Taiwan, 2010, pp. 1988–1993.
- [92] V. Raman, B. Xu, H. Kress-Gazit, Avoiding forgetfulness: Structured english specifications for high-level robot control with implicit memory, in: *Intl. Conf. on Intelligent Robots and Systems*, IEEE, Vilamoura, Algarve, Portugal, 2012, pp. 1233–1238.
- [93] L. Kunze, T. Roehm, M. Beetz, Towards semantic robot description languages, in: *Intl. Conf. on Robotics and Automation*, IEEE, Shanghai, China, 2011, pp. 5589–5595.
- [94] M. Tenorth, M. Beetz, KnowRob: A knowledge processing infrastructure for cognition-enabled robots, *Int. J. Robotics Res.* 32 (5) (2013) 566–590.

- 1692 [95] A. Miyazawa, P. Ribeiro, W. Li, A. Cavalcanti, J. Timmis, J. Woodcock, RoboChart: modelling and verification of the  
1693 functional behaviour of robotic applications, *Software & Systems Modeling* 18 (5) (2019) 3097–3149.
- 1694 [96] S. Schneider, *Concurrent and real-time systems: the CSP approach*, John Wiley & Sons, 1999.
- 1695 [97] F. Ciccozzi, D. D. Ruscio, I. Malavolta, P. Pelliccione, Adopting MDE for specifying and executing civilian missions of  
1696 mobile multi-robot systems, *IEEE Access* 4 (2016) 6451–6466.
- 1697 [98] I. Gavran, O. Mailahn, R. Müller, R. Peifer, D. Zufferey, Tool: accessible automated reasoning for human robot  
1698 collaboration, in: *Intl. Symp. on New Ideas, New Paradigms, and Reflections on Programming and Software*, 2018, pp.  
1699 44–56.
- 1700 [99] P. Detzner, T. Kirks, J. Jost, A novel task language for natural interaction in human-robot systems for warehouse logistics,  
1701 in: *Intl. Conf. on Computer Science & Education*, IEEE, Toronto, ON, Canada, 2019, pp. 725–730.
- 1702 [100] P. Forbrig, A. Bunea, Modelling the collaboration of a patient and an assisting humanoid robot during training tasks, in:  
1703 *Human-Computer Interaction*, Vol. 12182 of *Lecture Notes in Computer Science*, Springer, Copenhagen, Denmark, 2020,  
1704 pp. 592–602.
- 1705 [101] P. Forbrig, A. Dittmar, M. Kühn, A textual domain specific language for task models: Generating code for CoTaL, CTTE,  
1706 and HAMSTERS, in: *Symposium on Engineering Interactive Computing Systems*, ACM, Paris, France, 2018, pp. 5:1–5:6.
- 1707 [102] M. Webster, D. Western, D. Araiza-Illan, C. Dixon, K. Eder, M. Fisher, A. G. Pipe, A corroborative approach to  
1708 verification and validation of human–robot teams, *Intl. Journ. of Robotics Research* 39 (1) (2020) 73–99.

1709 **Appendix A. SHA Semantics**

1710 This appendix presents the semantics of SHA and shows, in Fig. A.15, an equivalent representation of  
 1711 the automaton of Fig. 3a that is fully compliant with the introduced syntax and semantics of SHA. The  
 1712 automaton of Fig. 3a must be understood as an equivalent, simplified representation of the automaton of  
 1713 Fig. A.15, where the latter differs from the former in the way the edge exiting location *cool* is connected to  
 1714 locations *high* and *low*.

1715 Complex systems constituted by multiple entities can be modeled as a combination of SHA, which form  
 1716 a **network**. To make  $n$  automata  $A_1, \dots, A_n$  (each one defined as in Definition 1) form a network, the  
 1717 following properties must be satisfied [21]. Every automaton  $A_i$  must be deterministic, i.e., there are no two  
 1718 (or more) edges, outgoing from a location of  $A_i$ , defined by the same event and whose edge conditions can be  
 1719 satisfied by the same valuation. Moreover, they must guarantee the following two semantic properties, called  
 1720 input-enabledness and time divergence. Let  $v'_{\text{var}} : W_i \rightarrow \mathbb{R}$  be a valuation for variables in  $W_i$ ,  $l_i \in L_i$  be a  
 1721 location of automaton  $A_i$  and let pair  $(l_i, v_{\text{var}})$  be a configuration of  $A_i$ . For every automaton  $A_i$ , and for all  
 1722 configurations  $(l_i, v_{\text{var}})$  and channel  $c \in C$ , there exists an edge  $c?$  that can be taken, i.e., the edge is enabled  
 1723 as the associated edge condition in  $\Gamma(W)$  is satisfied by  $v_{\text{var}}$ . Intuitively, this assumption ensures that every  
 1724 automaton can fire a transition in every possible configuration. Second, every automaton  $A_i$  always allows  
 1725 for executions such that if  $A_i$  is equipped with an extra clock which is never reset then, in all executions of  
 1726  $A_i$ , this clock cannot be bounded by any arbitrary integer constant (i.e., no Zeno executions are feasible).  
 1727 Finally, every automaton  $A_i$  is defined by considering the same set of channels ( $C_i = C_j$  when  $i \neq j$ ) and no  
 1728 pair of transitions, each one belonging to two different automata in the network, are built by referring to the  
 1729 same event  $c!$ . These properties are with no loss of generality. For example, disjointedness of channels can  
 1730 always be achieved by choosing properly defined symbol sets  $C_i$ . In addition, for simplicity, for the network  
 1731  $A_1, \dots, A_n$  to be composable the sets of real-valued variables must be pairwise disjoint ( $W_i \cap W_j = \emptyset$  when  
 1732  $i \neq j$ ). However, this constraint can be relaxed through a suitable extension of the semantics.

1733 In the following, we outline the semantics of an SHA network including  $n$  automata  $A_1, \dots, A_n$ , each  
 1734 defined as in Definition 1. The semantics of a composable network  $A_1, \dots, A_n$  is defined based on the  
 1735 configurations (of the network), each one being a tuple of the form  $(s_1, \dots, s_n)$  where every state  $s_i$  is a  
 1736 configuration of automaton  $A_i$ . There are two possible types of configuration changes realized, respectively,  
 1737 by *discrete transitions* and *time transitions*. A discrete transition occurs when one or more automata take an  
 1738 edge. In the latter case, at least two automata synchronize with each other. Synchronization among different  
 1739 automata inside a network occurs through the channels of set  $C$  [27]. Given a channel  $c \in C$  and two edges  
 1740 of two distinct automata, whose events are  $c!$  (the *sender*) and  $c?$  (the *receiver*), triggering an event through  
 1741 channel  $c$  causes both edges to fire simultaneously. Synchronization always requires at most one sender and  
 1742 possibly many receivers (even none). In Fig. 3b, the thermostat can trigger an event through channels *on!*  
 1743 and *off!* to start or stop heating the room. The triggered event is then received by the room automaton




Figure A.15: SHA modeling the room from the running example in Section 2 with detailed representation of how probability weights on receiving edges are handled.



1744 through labels `on?` and `off?`, which makes the corresponding edges fire. Taking an edge  $(l, c, \gamma, \xi, l')$  of  
1745 automaton  $A_i$  with configuration  $(l, v_{\text{var}})$  implies that the edge is enabled—i.e., all the conditions in  $\Gamma(W)$   
1746 associated with the edge are satisfied by the values defined by  $v_{\text{var}}$ . Upon taking the edge, the location of  
1747  $A_i$  changes from  $l$  to  $l'$  and the associated update  $\xi$  is executed, resulting in configuration  $(l', v'_{\text{var}})$ . Since  
1748 several automata may be involved in a synchronisation, and many updates can be executed simultaneously,  
1749 specific rules are needed to regulate their execution. The value of a variable  $w$  in  $v'_{\text{var}}$  is determined based on  
1750 the interpretation of  $w$ , i.e., whether  $w$  is a stochastic parameter or not. In the former case, upon entering a  
1751 location  $l' \in L_i$  such that  $\mathcal{D}_i(l')$  is defined, a realization of distribution  $\mathcal{D}_i(l')$  (e.g.,  $\mathcal{N}(\mu_H, \sigma_H^2)$  in Fig. A.15)  
1752 defines the value of  $w$  in  $v'_{\text{var}}$  (e.g.,  $\theta$  in Fig. A.15); otherwise, when  $\mathcal{D}_i(l')$  is not defined, the value of  $w$  in  
1753  $v'_{\text{var}}$  and in  $v_{\text{var}}$  is the same [17]. In the latter case,  $w$  is not interpreted as a randomly distributed parameter  
1754 and its value in  $v'_{\text{var}}$  is the value of the assignment associated with  $w'$  in  $\xi$ , that is obtained by evaluating  
1755 every non-primed variables of the constraint with values from  $v_{\text{var}}$ . The configuration  $(l', v'_{\text{var}})$  is such that  
1756 the valuation  $v'_{\text{var}}$  satisfies the invariant  $\mathcal{I}_i(l')$ .

1757 Besides randomly distributed variables, in SHA, probability measures can be associated with delays to  
1758 model the elapsing of time in the network, hence the wait between the occurrence of two discrete transitions.  
1759 According to [21], the adopted probabilistic semantics is based on the “principle of independence” among  
1760 automata in the network. Upon the firing of an edge, for every automaton  $A_i$  in the network, a delay  
1761  $d_i$  models the time  $A_i$  waits before taking an edge for event  $c!$ , for some  $c \in C$ . If no edges for event  
1762  $c!$  originate from  $l'$ , then  $d_i$  is  $\infty$ . Otherwise, let  $d_{\min}(l', v'_{\text{var}})$  be the minimum delay that automaton  $A_i$   
1763 should wait before an edge whose event is  $c!$ , and departing from  $l'$ , is enabled; and let  $d_{\max}(l', v'_{\text{var}})$  be  
1764 the maximum delay that automaton  $A_i$  can wait before all edges, for events  $c!$ , with  $c \in C$ , exiting  $l'$  are  
1765 disabled (note that both values are a function of the invariant  $\mathcal{I}_i(l')$ , of the edge conditions and of the  
1766 current valuation  $v'_{\text{var}}$ ). If  $d_{\min}(l', v'_{\text{var}})$  is not defined, then  $d_i$  is  $\infty$ ; otherwise, if  $d_{\min}(l', v'_{\text{var}})$  is defined,  
1767  $d_i$  is a realization of the probability distribution  $\mu_i(l', v'_{\text{var}})$ . If  $d_{\max}(l', v'_{\text{var}})$  is finite, then  $\mu_i(l', v'_{\text{var}})$  is a  
1768 uniform distribution over the interval  $[d_{\min}(l', v'_{\text{var}}), d_{\max}(l', v'_{\text{var}})]$ ; otherwise,  $\mu_i(l', v'_{\text{var}})$  is an exponential  
1769 distribution over  $[d_{\min}(l', v'_{\text{var}}), \infty)$ . If  $d_i$  is  $\infty$ , by input-enabledness, then  $A_i$  can take an edge, whose event  
1770 is  $c?$ , for some  $c \in C$ . Otherwise, by definition of  $d_i$ , after  $d_i$  time units from the current discrete transition,  
1771 automaton  $A_i$  can surely take an edge, whose event is possibly  $c!$ , for some  $c \in C$ . Since the network consists  
1772 of  $n$  automata, the minimum allowed progress  $d_m$  is selected among the  $n$  delays  $d_1, \dots, d_n$ . If  $d_m$  is finite,  
1773 then  $d_m$  is the time the network waits before an automaton performs a new discrete transition.

1774 The wait between the execution of two discrete transitions, lasting a generic  $\delta > 0$  time units, is a timed  
1775 transition, i.e., a configuration change such that no location of the automata in the network is modified but  
1776 values of the variables evolve because of the elapsing of time. The configuration of the  $i$ -th automaton at the  
1777 end of this wait is  $(l'', v''_{\text{var}})$ , with  $l'' = l'$ , where  $(l', v'_{\text{var}})$  is the configuration whence the timed transition  
1778 starts. All the variables of the set  $W_i$  evolve according to the flow conditions  $\mathcal{F}_i(l')$ . In the case of clocks  
1779  $x \in X_i$ , for instance, they are incremented by the value  $\delta$ , hence,  $v''_{\text{var}}(x) = v'_{\text{var}}(x) + \delta$  holds. The value of  
1780 the other variables is determined based on the differential equation specified by  $\mathcal{F}_i(l')$ . With the adopted  
1781 semantics,  $\delta$  is the value  $d_m$  calculated at the occurrence of the last discrete transition.

1782 At the end of the time interval lasting  $d_m$  units of time, the automaton  $A_i$  such that  $d_i = d_m$  holds  
1783 performs a discrete transition for some event  $c!$ , with  $c \in C$ . If several edges are enabled in  $(l', v'_{\text{var}})$ ,  
1784 probability distribution  $\mathcal{P}(l')(c!, \gamma', \xi', l'') \in [0, 1]$  with  $l'' \in L_i$ ,  $\gamma' \in \Gamma_i(W_i)$ , and  $\xi' \in \wp(\Xi_i(W_i))$  determines  
1785 how likely the system is to evolve in one direction rather than the other. In Fig. A.15,  $p_L$  and  $p_H$  are the  
1786 probability of the switching of the heating when the window is open or closed, respectively, which takes place  
1787 after the synchronization between the two automata has been achieved through channel `on`. Channels `onH`  
1788 and `onL` in Fig. A.15 model a probabilistic choice and are not intended for synchronizing the two automata.

1789 We remark that some of the models presented in this work do not conform with the disjointness of the  
1790 set of real-valued variables, hence two or more automata can use the same variable. This, however, is with  
1791 no loss of generality in our work, because it is always possible to introduce suitable transitions and local  
1792 copies of the shared variables and build a network such that all sets of real-valued variables are pairwise  
1793 disjoint. An automaton  $\mathcal{A}_1$  can always make an automaton  $\mathcal{A}_2$  change a variable  $v$  in  $W_2$  by means of two  
1794 synchronizing edges with a dedicated event, possibly representing the operation to be carried out on  $v$ .

## 1795 Appendix B. DSL for Framework Validation Scenarios

1796 This Appendix contains the DSL configuration of scenarios DPa, DPb, and DPc. The complete .dsl file is  
1797 constituted by the concatenation of Listings 5 through 9.

Listing 5: DSL section defining layout areas (i.e., the rectangles highlighted in Fig. 12b) and POIs: specifically, the entrances to the three offices, to the waiting room and emergency room, the two cupboards, main entrance, and robot's recharge station. As all scenarios are set in the same layout, the DSL features only one layout definition.

```
1798 1 param measurement_unit cm
1799 2 define layout:
1800 3     area a1 in (0.0,110.0) (1550.0,299.5)
1801 4     area a2 in (0.0,110.0) (185.0,850.0)
1802 5     area a3 in (0.0,672.5) (1550.0,850.0)
1803 6     area a4 in (1352.0,110.0) (1550.0,850.0)
1804 7     area a5 in (2970.0,110.0) (4512.5,299.5)
1805 8     area a6 in (2970.0,110.0) (3155.0,850.0)
1806 9     area a7 in (2970.0,672.5) (4512.5,850.0)
1807 0     area a8 in (4322.0,110.0) (4512.5,850.0)
1808 1     area a9 in (1945.0,0.0) (2670.0,695.0)
1809 2     area a10 in (1352.0,110.0) (3155.0,425.0)
1810 3
1811 4     poi OFF1 in (200.0, 200.0)
1812 5     poi OFF2 in (4400.0, 200.0)
1813 6     poi OFF3 in (4400.0, 700.0)
1814 7     poi R1a in (1200.0, 680.0)
1815 8     poi R1b in (400.0, 270.0)
1816 9     poi R2 in (4000.0, 270.0)
1817 0     poi CUP1 in (1400.0, 450.0)
1818 1     poi CUP2 in (3000.0, 450.0)
1819 2     poi ENTR in (2300.0, 600.0)
1820 3     poi RECH in (4250.0, 450.0)
```

Listing 6: DSL section defining robot Tbot and its features. As illustrated in Section 6, it is a TurtleBot3 Waffle Pi starting with 90% of charge.

```
1821 1 define robots:
1822 2     robot Tbot in (2300.0, 400.0) id 1 type turtlebot3_wafflepi charge 90
```

Listing 7: DSL section defining the human subjects and their features. Patients (P1a, P1b, P1c, and P2c) all have sick fatigue profiles, and only P2c belongs to the elderly age group. Doctors (D1a, D1b, D1c, and D2c) all have healthy fatigue profiles and belong to the elderly age group, except for D2c. Walking speeds are set to 40cm/s for patients, and 100cm/s for doctors.

```
1823 1 define humans:
1824 2     human P1a in (2300.0, 600.0) id 1 speed 40.0 is young_sick freewill
1825 3     normal
1826 4     human D1a in (4400.0, 700.0) id 2 speed 100.0 is elderly_healthy freewill
1827 5     low
1828 6
1829 7     human P1b in (2300.0, 600.0) id 1 speed 40.0 is young_sick freewill
1830 8     normal
1831 9     human D1b in (4400.0, 700.0) id 2 speed 100.0 is elderly_healthy freewill
1832 0     normal
1833 1
1834 2     human P1c in (2290.0, 600.0) id 1 speed 40.0 is young_sick freewill high
```

```

18359     human P2c in (2400.0, 580.0) id 2 speed 40.0 is elderly_sick freewill
1836     normal
18370     human D1c in (200.0, 200.0) id 3 speed 100.0 is elderly_healthy freewill
1838     low
18391     human D2c in (4400.0, 700.0) id 4 speed 100.0 is young_healthy freewill
1840     normal

```

Listing 8: DSL section defining the service sequences (i.e., the robotic missions). As described in Section 6.2, each scenario corresponds to a mission declaration. Service sequences are defined as in Table 6 and Table 8.

```

18411  define mission DPa:
18422      do robot_leader for P1a with target R1b
18433      do robot_follower for D1a with target CUP1
18444      do robot_follower for D1a with target R2
18455      do robot_leader for P1a with target R2
18466
18477  define mission DPb:
18488      do robot_leader for P1b with target R1a
18499      do robot_transporter for D1b with target CUP2
18500      do robot_follower for D1b with target R2
18511      do robot_leader for P1b with target R2
18522
18533  define mission DPc:
18544      do robot_leader for P1c with target R1a
18555      do robot_leader for P2c with target R2
18566      do robot_transporter for D1c with target CUP1
18577      do robot_follower for D2c with target CUP2
18588      do robot_follower for D2c with target OFF3
18599      do robot_leader for P1c with target OFF1
18600      do robot_leader for P2c with target OFF3
18621
18622  define mission R-DPa:
18623      do robot_leader for P1a with target R2
18624      do robot_follower for D1a with target CUP1
18625      do robot_follower for D1a with target R2
18626
18627  define mission R-DPb:
18628      do robot_transporter for D1b with target CUP2
18629      do robot_follower for D1b with target R2
18630      do robot_leader for P1b with target R2
18631
18632  define mission R-DPc:
18633      do robot_leader for P2c with target R1b
18634      do robot_leader for P1c with target R1a
18635      do robot_transporter for D1c with target CUP1
18636      do robot_follower for D2c with target CUP2
18637      do robot_follower for D2c with target OFF3
18638      do robot_leader for P1c with target OFF1
18639      do robot_leader for P2c with target OFF3

```

Listing 9: DSL section defining the queries to be performed for the design-time analysis. Queries defined in this Listing yield the results shown in Table 7 and Table 9.

```

18801 define queries of mission DPa:
18812     compute probability_of_success with duration 400 runs auto
18823     compute probability_of_success with duration 350 runs auto
18834     compute probability_of_success with duration 300 runs auto
18845     compute expected_charge with duration 400 runs auto
18856     compute expected_fatigue with duration 400 runs auto
18867
18878 define queries of mission DPb:
18889     compute probability_of_success with duration 520 runs auto
18890     compute probability_of_success with duration 450 runs auto
18901     compute probability_of_success with duration 400 runs auto
18912     compute expected_charge with duration 520 runs auto
18923     compute expected_fatigue with duration 520 runs auto
18934
18945 define queries of mission DPc:
18956     compute probability_of_success with duration 1500 runs auto
18967     compute probability_of_success with duration 1400 runs auto
18978     compute probability_of_success with duration 1300 runs auto
18989     compute expected_charge with duration 1500 runs auto
18990     compute expected_fatigue with duration 1500 runs auto
19001
19022 define queries of mission R-DPa:
19023     compute probability_of_success with duration 300 runs auto
19024     compute probability_of_success with duration 250 runs auto
19025     compute probability_of_success with duration 200 runs auto
19026     compute expected_charge with duration 300 runs auto
19027     compute expected_fatigue with duration 300 runs auto
19028
19029 define queries of mission R-DPb:
19030     compute probability_of_success with duration 350 runs auto
19031     compute probability_of_success with duration 320 runs auto
19032     compute probability_of_success with duration 300 runs auto
19033     compute expected_charge with duration 350 runs auto
19034     compute expected_fatigue with duration 350 runs auto
19035
19036 define queries of mission R-DPc:
19037     compute probability_of_success with duration 1500 runs auto
19038     compute probability_of_success with duration 1400 runs auto
19039     compute probability_of_success with duration 1300 runs auto
19040     compute expected_charge with duration 1500 runs auto
19041     compute expected_fatigue with duration 1500 runs auto

```