



Shared Task Representation for Human–Robot Collaborative Navigation: The Collaborative Search Case

Marc Dalmaso^{1,2} · J. E. Domínguez-Vidal^{1,2} · Iván J. Torres-Rodríguez^{1,2} · Pablo Jiménez¹ · Anaís Garrell^{1,2} · Alberto Sanfeliu^{1,2}

Accepted: 25 September 2023
© The Author(s) 2023

Abstract

Recent research in Human Robot Collaboration (HRC) has spread and specialised in many sub-fields. Many show considerable advances, but the human–robot collaborative navigation (HRCN) field seems to be stuck focusing on implicit collaboration settings, on hypothetical or simulated task allocation problems, on shared autonomy or on having the human as a manager. This work takes a step forward by presenting an end-to-end system capable of handling real-world human–robot collaborative navigation tasks. This system makes use of the Social Reward Sources model (SRS), a knowledge representation to simultaneously tackle task allocation and path planning, proposes a multi-agent Monte Carlo Tree Search (MCTS) planner for human–robot teams, presents the collaborative search as a testbed for HRCN and studies the usage of smartphones for communication in this setting. The detailed experiments prove the viability of the approach, explore collaboration roles adopted by the human–robot team and test the acceptability and utility of different communication interface designs.

Keywords Human–robot collaboration · Human–robot collaborative navigation · Human–robot interaction · Multi-agent planning · Motion planning · Task representation · Object search

1 Introduction

Some tasks where humans collaborate with each other entail a displacement across a given scenario. Moreover, such displacement is a fundamental part of the task and requires to be coordinated between the collaborating actors (e.g. ware-

house logistics, delivery, sports, transportation through joint manipulation...). This is what we call Human Collaborative Navigation, which is frequently entangled with manipulative tasks or other actions.

Human–Robot Collaboration (HRC) has made many advances focusing on constrained tasks. In fact, some specific settings have received so much interest from the industry that they are sometimes treated as the actual definition of HRC (e.g. industrial shared workspaces [1–3]). Nonetheless, HRC embraces a wide range of challenges and, as in human collaboration, an important subset of them involve navigation. These challenges inspire a sub-field of HRC research known as Human–Robot Collaborative Navigation (HRCN).

Until now Human–Robot Collaborative Navigation (HRCN) lacked a representative constrained testbed, a movement-only real-world collaborative navigation task recognised by its solvers and explicitly plannifiable. Therefore, real-world applications have focussed on improving performance in specific cases like shared autonomy, human control over robot teams or implicit collaboration scenarios (e.g. person-following, side-by-side navigation, handover with motion). In a more general sense, HRCN should include multitasking, human–robot joint-tasks, multifinal goals and

✉ Marc Dalmaso
marc.dalmaso@upc.edu

J. E. Domínguez-Vidal
jdominguez@iri.upc.edu

Iván J. Torres-Rodríguez
ivan.jesus.torres@upc.edu

Pablo Jiménez
pjimenez@iri.upc.edu

Anaís Garrell
anaís.garrell@upc.edu

Alberto Sanfeliu
alberto.sanfeliu@upc.edu

¹ Institut de Robòtica i Informàtica Industrial (CSIC-UPC), Llorens i Artigas 4-6, 08028 Barcelona, Spain

² Universitat Politècnica de Catalunya - BarcelonaTech (UPC), Jordi Girona, 31, 08034 Barcelona, Spain

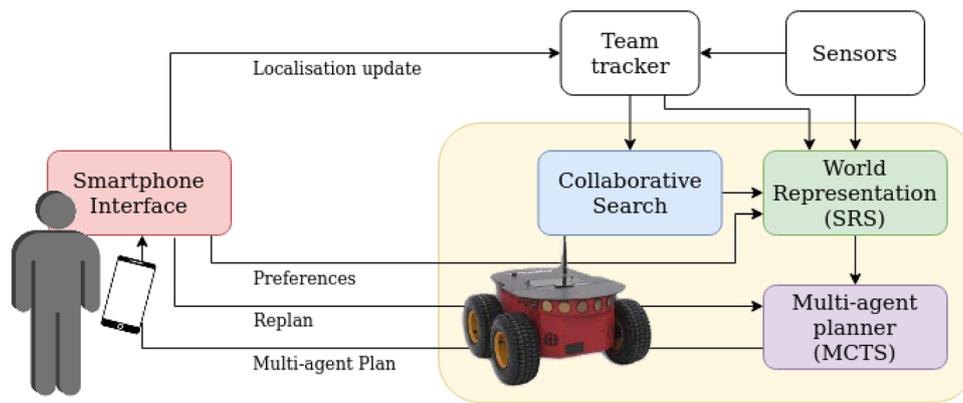


Fig. 1 Main contributions. We built a system capable of handling generic collaborative navigation tasks. The main contributions are building a knowledge representation model (green), designing a planner for heterogeneous human–robot teams (purple) and defining a bench-

mark for human–robot collaborative navigation tasks (blue). Finally, we developed an interface to handle the team interaction (red). Acronyms used in the figure correspond to Social Reward Sources (SRS) and Monte–Carlo Tree Search (MCTS). (Color figure online)

actions, task allocation, team–role assignment, shared knowledge representation, mental models and agent preferences. Consequently, a system tackling HRCN should be able to represent the previous and include the necessary tools to allow communication, model reconciliation processes, role assignment, planning and negotiation.

The overall goal of this work is to build a system capable of handling a generic collaborative navigation task. This implies building a knowledge representation model, designing a planner for heterogeneous human–robot teams, defining a benchmark for human–robot collaborative tasks and developing an interface to handle interaction and multitasking (Fig. 1).

To test the HRCN system, we propose the human–robot collaborative search. In this context, coordination means agreeing in a plan to explore the search space in a complementary fashion (i.e. avoiding exploring the same or concurrent areas) and achieve the collaborative goal of finding one object. Searchers keep track of the whole task progress while observing teammates’ behaviour, thus inferring others’ future actions and planning exploration strategies.

The present document is organised as follows. Section 2 introduces the related work. Section 3 explains the Social Reward Sources (SRS) model, the knowledge representation model, whilst Sect. 4 develops the multi-agent planner. Section 5 defines the collaborative search benchmark and 6 discusses the emerging team–roles from the proposed collaboration. Afterwards, Sect. 7 describes the app interface used to enable human–robot communication. The conducted experiments and their results are explained in Sect. 8. Finally, a summary discussion of the overall work and the extracted conclusions are presented in Sects. 9 and 10, respectively.

2 Related Work

In this section we discuss the work on previous articles to depict the state of the art in this field. We review task representation approaches for human–robot teams and current approaches to human–robot collaborative navigation settings. Additionally, we discuss previous approaches human–robot communication interfaces that inspired our application.

Section 2.1 reviews different approaches to shared task representation in human–robot teams. Section 2.2 discusses current approaches and benchmarks in HRCN and, finally, Section 2.3 describes significant references in the development of human–robot interfaces.

2.1 Shared Task Representation in Human–Robot Teams

We understand a shared task representation as the compound of all information related to a task held in common by all team members. Such construction may be achieved at different levels, from simple assumptions of commitment of the teammates to full understanding of their knowledge, decision making processes and goals. We consider it to be useful once it enables meaningful team communication that changes the task execution. Bearing that in mind, here we review different approaches to knowledge representation and human–robot communication that might build a bridge between human and robot world understanding.

On the one hand, a core challenge is human understanding of the robots’ capabilities, preferences, plans and goals. Chakraborti et al. [4] discusses the meaning of different con-

cepts arising from this endeavor (e.g. legibility, transparency, explicability, predictability) and its counterparts (e.g. dissimulation, obfuscation), all of them concerning either goals or plans. Their review mainly focuses on path and motion planning, mostly relying in humans' learned physical intuition to enable the information transfer process. Similarly, other works establish human–robot communication through the exertion of forces [5, 6] or commonly known physical interfaces (e.g. joysticks [7]). One alternative to convey higher level information to the human in human–robot teams is the usage of augmented reality (AR) [3], although it burdens the human with the usage of an interface it helps to convey spatial information that might be difficult to express in other means.

On the other hand, multiple fields work on robot's understanding of the human capabilities, preferences, plans and goals. For example, recently Tuli et al. [8] built an ontological-based system for human intention inference in assembly operations (as in current goal), Liu et al. [1] claimed to improve task scheduling in shared workspace settings through introducing dynamic and stochastic representations of the human task performance model and Rudenko et al. [9] surveys advances in pedestrian trajectory prediction. Architecturally speaking, classically the human was treated as an element of the environment, but recently the paradigm of considering human–robot collaboration using a multi-agent view is receiving more attention [10]. Building a multi-agent plan is of special importance for human–robot teams tackling collaborative navigation tasks. These tasks ask for longer planning time horizons than other settings, and human teammates' contribution can completely change the usefulness of a given robot plan. Some ontological approaches to motion planning can store and receive profound knowledge of the task (e.g. Tuli et al. [8]), but they are usually implemented in fully known environments. Moreover, they are frequently used in architectures where task allocation processes are disentangled from motion planning algorithms (e.g. [11]). This disentanglement can prove to be problematic in dynamic settings where tasks time and cost may experience major changes, such as in social navigation and HRCN. That being said they are ubiquitous in systems dealing with temporal constraints such as time window, synchronisation or precedence constraints (e.g. Nunes et al. [12]), as there seem to be no motion planning methods that deal with them. Similarly, multi-agent systems dealing with multitask systems should account for any arising interdependence [13].

Interestingly, Elliot & Thrash [14] work supports the simplification of representing human goals as approach and avoidance motives. Some human behaviour modelling approaches like the different variations of the social force model [15, 16] also follow this trend. This is of special interest, as many robot navigation and learning methods build upon the concept of rewards, positive or negative feedbacks received upon taking some actions or achieving some states.

This opens up a communication opportunity, as human interpretation of robot actions in terms of action goals and means to achieve them is similar to how other human actions are conceived [17]. More so, Lin et al. [18] suggest that complex rewards from learned behaviours may be disentangled in sub-rewards with semantic significance. Similarly, cost partitioning approaches [19] suggest that multiple rewards representing different features may be combined in an additive fashion. We believe a modular representation of the sources of reward present in the task world can provide a flexible task representation, capable of adapting to changes in the environment and the actors. Building from these, in this work we propose to build a shared task representation for HRCN from a modular reward definition architecture capable of representing spatial properties and temporal constraints.

2.2 Human Robot Collaborative Navigation

Currently, there isn't a clear definition of what is and what isn't human–robot collaboration in navigation tasks. Some authors argue that pedestrian avoidance is a collaborative or cooperative task due to the altruistic behaviours that arise in these situations [20–22]. Here, we will not generally review social navigation approaches, but works handling human–robot navigation settings where they share a common goal.

Two of the most studied collaborative scenarios in navigation tasks are the shared autonomy and human manager settings. Most publications concerning these pertain to the search and rescue research field [7, 23, 24]. We may label them as examples of *telematic HRCN*, as in both scenarios the human isn't present in the navigation task space and only contributes to the shared task through either controlling the robot actions or managing team members coordination. Some of these works consider the presence of human team members in the scene, but generally their interaction with the robots is handled indirectly through the manager.

Some commonly studied scenarios where humans and robots share a workspace and perform shared navigation tasks are person-following [25], handover with motion [26] and side-by-side navigation [27–29]. Though in these scenarios there exist a shared task, the problem statement generally considers a human performing the collaboration subconsciously. Actually, achieving natural behaviour (their usual behaviour) in humans who are over-conscious of the task can be difficult in the experiment design of such collaborative tasks. Assuming subconscious collaboration, we will label such tasks as *implicit HRCN* scenarios.

Alternatively, Liu et al. [6] presents a shared manipulation system for human–robot transportation. Conversely, in this setting the human is consciously engaging in the collaborative task. The proposed approach, however, focuses on the robot's dynamic control and the human is only treated as an environmental payload input. Thus, we consider the robot

as unaware of the interaction and also label the proposed solution as *implicit HRCN*.

We would label a collaborative scenario as *explicit HRCN* whenever all agents are aware of the task. Arguably, only in explicit HRCN there may be human–robot teams. However, there seems to be a lack of research tackling this kind of tasks.

2.3 Human–Robot Interfaces

The interface used to enable human–robot communication in this work is based on the previous work from Kohler et al. [30], which used mobile phones using Android OS in order to communicate with robots using the ROS framework [31]. It is also based on Fogg’s previous work [32] on persuasive design. More specifically, on their concept of trigger. Our interface relies heavily on visual cues drawn upon a representation of the world. Similar approaches are frequently used [3, 13] as they can convey spatial information that would be otherwise difficult to express.

3 Social Reward Sources

The Social Reward Sources (SRS) model is a world and task representation model constructed using sources of reward as building blocks. It is a subjective representation where meaningful entities to the agents (robot or human) are modelled as sources of reward or *reward sources*. Therefore, the SRS model is constructed as a set of reward sources $\psi \in \Psi$.

Important preceding research on social navigation involves the modelling of tasks using virtual forces (Social Force Model (SFM)) [15, 16]. Our model inherits its advantages as it originates from rethinking the late tendency to use said forces as costs in the planning process [27, 28]. Also, it generalises reward generation to avoid previous ad-hoc solutions, adapts to complex semantic space definitions [33, 34] and models the existence of shared goals, multitasking and multifinal actions. Additionally, it incorporates needed dependency relations between tasks [13]. Another advantage is that humans find it intuitive to generate plans using this model (see Sect. 3.3.1).

Solutions to task representation in human–robot settings in the literature are usually task-focused [6–8, 10] and do not generalise to other scenarios. For instance, Liu et al. [6] task representation is done through a control loop and can only be applied to collaborative transportation, previous SFM applications [15, 16, 27] can only be applied to the target task they are handling and task representation in multi-robot collaborative systems [35] is not designed to be used for human–robot communication. However, although we can’t directly compare to them, it is worth noting that this model has shown potential in human–robot shared world representation of navigation tasks [36].

Section 3.1 defines the concept of source of reward ψ and its properties. Later, Sect. 3.2 introduces the world representation constructed through the interaction of many sources. Section 3.3 discusses the usage of the SRS model in HRI, while in 3.4 the multi-agent extension of the model is presented.

3.1 Source of Reward

Ultimately, a social reward source $\psi \in \Psi$ is a generative model that defines a reward function along all the search space $r(\psi) = f(x, y)$ (this work focuses in \mathbb{R}^2 navigation, but could be applied to \mathbb{R}^3 or a robot joint space, for example). Reward sources act as building blocks of a subjective world representation. We define it as subjective because each source of reward may provide a different feedback to different agents, even under the same conditions. For example, rewards may vary due to risk, capabilities, effort, personal preferences or eagerness to do a task or to follow someone’s instructions. Hence, the constructed world representation is subjective to a particular agent, from now on a_0 .

We want to use SRSs to build a unified subjective world representation of all external events Ψ_W (Fig. 2). This includes perceived and interpreted environment influences E (from obstacles avoidance to conceptual abstractions such as areas in a sports field), the pleasantness, cost or effort directly related to performing certain actions A , rewarding feedback due to progressing in tasks in hand or reaching goals T and feedback emerging from the interaction with other agents I (considering human or robot sharing any task with a_0).

3.1.1 Source Definition

As introduced before, sources of reward may have non-trivial spatial definitions. Sources of reward may exist attached to physical objects, e.g. obstacle avoidance rewards. Others may be defined over human constructed concepts, such as the previously mentioned areas in a sport field. For practicality, even classic navigation goals should contain some specified tol-

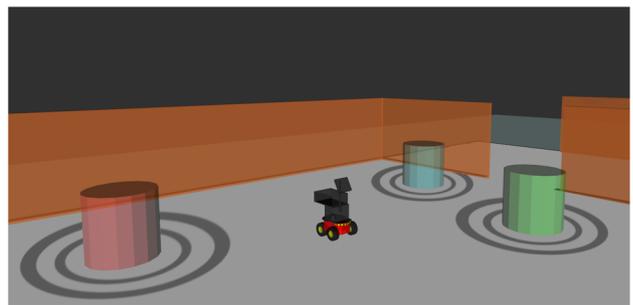


Fig. 2 World representation. In the Social Reward Sources model, all entities and concepts influencing the actions of the robot are modelled as sources of reward

erance, which consequently defines an accepted zone. One feature we deemed to be common in most cases is the existence of a boundary. Lots of sources have a binary effect, showing clearly opposed effect whether a_0 is located inside or outside the boundary (e.g. pressing a button). Others have a similar effect while including some transition (e.g. obstacle avoidance or positioning oneself in a cinema).

We consider that all sources ψ have a boundary β . This boundary is a closed surface in the search space (in \mathbb{R}^2 the boundary would enclose a determined area). The reward production (reward/t) of the source is defined by two functions: r_i and r_e , corresponding to internal and external reward ratio functions:

$$r(p_a) \begin{cases} r_i(\vec{p}_\beta^a), & \text{if } p_a \subseteq \bar{\beta} \\ r_e(\vec{p}_\beta^a), & \text{otherwise} \end{cases} \quad (1)$$

where p_a is the agent location, \vec{p}_β^a is a vector pointing p_a from its nearest boundary point (Fig. 3) and $\bar{\beta}$ is the closure of the boundary surface (all points in the interior of the boundary or on the boundary). The total reward output capability of the source for a given location p is as follows:

$$r(p, \Delta t_i) = \int_{t_{i0}}^{t_{if}} r(p) \cdot \delta t \quad (2)$$

where,

$$\Delta t_i = t_{if} - t_{i0} \quad (3)$$

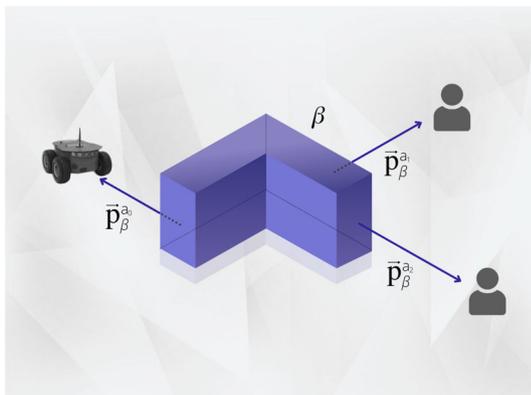


Fig. 3 Reward definition. In front of a source, the reward perceived by each agent is defined as a function of the shortest vector between its boundary and their location \vec{p}_β^a . This vector is, by definition, perpendicular to the boundary surface β . Likewise, its dimensionality will be equivalent to that of the source definition

3.1.2 Reward Consumption

Some sources of reward should be continuously outputting reward. As a clear example, a building containing fire smoke should be avoided at all times (except if the agent is a fire-fighter, another nice example for the subjectivity of reward sources). Other rewards, however, can only be collected once or a limited number of times. This is a dynamic usually observable in positive rewards. Sampling methods classically only use costs (negative rewards), as positive rewards generate sinks (negative costs). Even in q-learning positive rewards are usually only used in situations where they may not be retaken (goal states, accumulating points...).

Ultimately, in all applications using rewards the world representation system should keep track of the obtained rewards. This is usually solved ad-hoc, as the research focus is normally targeting the planning or decision making method. The SRS model shifts this responsibility to each source implementation, aiming to add flexibility in dealing with open world problems.

Each source ψ has a determined reward production capability R . The final reward produced by the source is as follows:

$$r_f(p, \Delta t_i) \begin{cases} r(p, \Delta t_i), & \text{if } r(p, \Delta t_i) \leq R_t \\ R_t, & \text{if } r(p, \Delta t_i) > R_t \end{cases} \quad (4)$$

where,

$$R_{t+1} = R_t - \sum^a r(p_a, \Delta t) \quad (5)$$

$$R_0 = R \quad (6)$$

The available reward may be kept constant if the source is defined with $R \rightarrow \infty$ or may be consumed nearly instantly if $R = \max_p r$ (Fig. 4). We may visualize sources as wells, but outputting reward from a reward pool instead of water.

3.1.3 Environment Rewards

Rewards linked to elements of the environment. The star example of these rewards is building obstacle avoidance, which may be achieved through an exponential model (as in the SFM [15]). Changing the environment changes their rewards, but should not affect any of the following reward sets.

$$\psi \in \Psi_E$$

3.1.4 Action Rewards

Action rewards are rewards linked to the actions of the agent. They can be seen as an internal representation of agent a_0 .

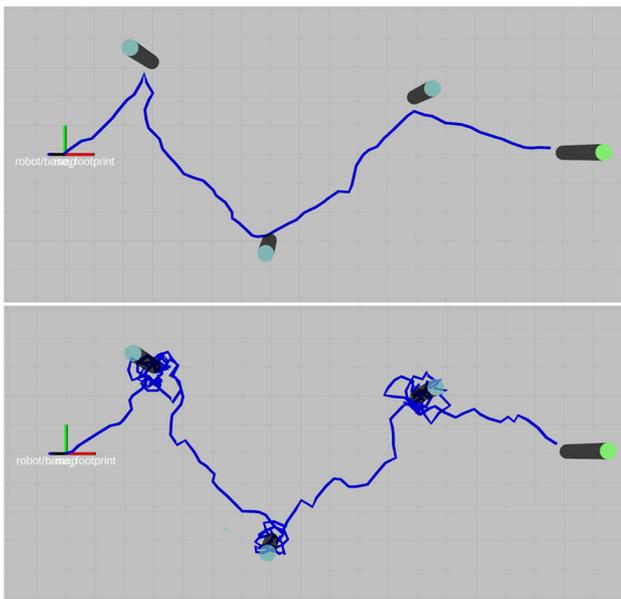


Fig. 4 Reward consumption. Even with a similar distribution of positive reward sources, varying the available reward R can generate significantly different behaviours. These images showcase the obtained path using RRT^* in two different settings. The blue and green sources have reward productions r_b and r_g , correspondingly. The blue sources on the above one are constructed with $R = \max_p r_b$, while the ones below have $R \rightarrow \infty$. In both cases, $r_g \gg r_b$ and the green source should be collected the last (see precedence relationships in Sect. 3.2.1). The green source is added to aid visualization, forcing the plan of the second scenario to move towards it (see objective function details in Sect. 4.5). (Color figure online)

If it is a robot, it may represent energy cost and be defined over robot configuration. If a_0 is a person, it may be tuned to represent fatigue, motivation or comfort. These rewards can have a major role over task selection outcomes. Specially in the task assignment processes of multi-agent planning.

Note that action sources are usually defined in a different space than the rest. For instance, action sources used in the presented experiments are defined in a one-dimensional space whose only axis is the agent velocity. Alternatively, other settings focusing in comfort may include variables defining the agent's body or arm configuration.

$$\psi \in \Psi_A$$

3.1.5 Task Rewards

Task rewards are linked to a_0 's goals and objectives. This includes sources of reward whose existence is subordinated to the tasks at hand. This set, however, is restricted only to the sources whose boundary (including its position) is independent of all agents' actions.

When task planning over the SRS model, one or various tasks with different goals may coexist in the same world representation. In conjunction with Action rewards,

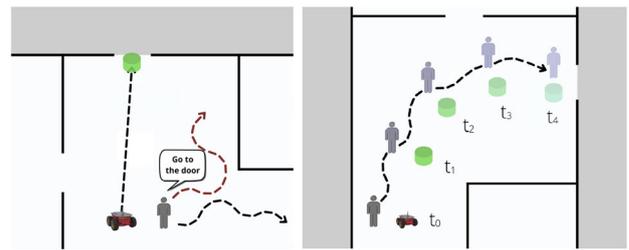


Fig. 5 Task versus interaction. The main difference between task and interaction sources of rewards is their spatial definition. Sources whose location is independent of any team member's movement actions are considered task sources, even if they are generated through direct interaction (e.g. the verbal command stated in the left image: "go to the door"). Interaction rewards are those whose spatial definition is generated in one of the team members' reference. Typical examples would be side-by-side navigation, which is depicted in the right image. Notice the goal moves according to the agent's actions and it would change if they took another path

this approach unifies task selection and navigation planning. It enables balancing tasks' rewards while taking into account well-grounded movement cost estimations. Furthermore, new requirements or tasks may be added or deleted on execution.

Notice that suggesting another agent to go to a certain fix location would fall into this reward source category. Even though the reward is generated due to the interaction with another agent, the task itself is independent of that agent position (Fig. 5).

$$\psi \in \Psi_T$$

3.1.6 Agent Interaction Rewards

Sources of reward related to tasks directly involving spatial interaction with other agents may be classified as I . Their separate treatment comes up due to their spatial definition, as it is referenced relative to the agents' position. For example, goals handling side-by-side navigation would fall into this category, as they are directly dependent on the companion's position and orientation. This example, and the one presented in the previous section, may be clarified when considering task progress in the future. When predicting or planning future states, handling goals dependent on other agents actions asks for a specific treatment (Fig. 5).

$$\psi \in \Psi_I$$

3.2 SRS World Representation

A SRS task world representation W_{SRS} is constructed by a set of sources of reward $\Psi = \{\psi_0, \psi_1 \dots \psi_n\}$. A complete world representation will be identified as Ψ_W and it may be subdivided in sets of each of the previous categories.

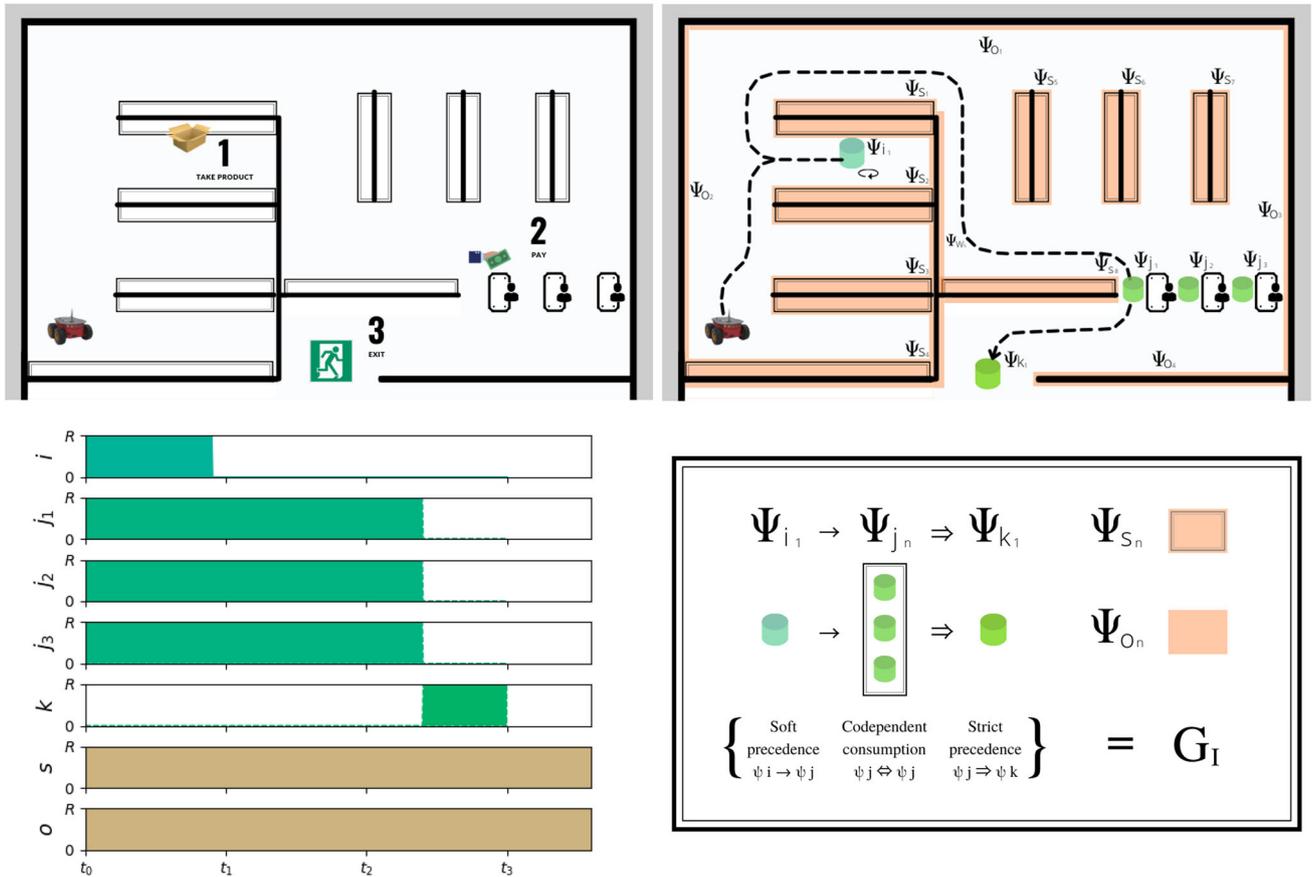


Fig. 6 Supermarket example. From top to bottom, left to right: **a** Toy example depicting a robot buying in a modern day supermarket. **b** Depiction of the relevant sources attached to the current scenario and

a possible plan to satisfy the task at hand. **c** Time graph depicting the available reward on each source along the execution the plan presented in **b**. **d** Supermarket sources' interdependence

$$\Psi_W = \{\Psi_E, \Psi_A, \Psi_T, \Psi_I\} \tag{7}$$

As per definition, Ψ_W represents a multi-objective setting. Classic methods would separately tackle task selection and planning problems, but here we propose a model where planners should confront both simultaneously. Notice that this framework naturally takes into account the existence of multi-final actions, as reward from different sources may be collected simultaneously by the agents.

3.2.1 Source Interdependence

Reward sources in a world representation can present interdependence. These relations should be encoded with the source set, as our world representation should be usable for planning (Fig. 6). As of now, we defined three source interdependence relations.

Strict precedence ($\psi_i \Rightarrow \psi_j$): Some sources may only be rewarding after consuming another source's reward (e.g. going to a closed door may only be rewarding once we collected the key). □

Soft precedence ($\psi_j \rightarrow \psi_i$): Some sources may only be rewarding before consuming another source's reward, though their consumption is not mandatory for that to be consumed (e.g. shifting one's path from to pass through certain mid-points may only be rewarding until reaching our goal, such as prioritising shadows in summer sunny days). □

Codependent consumption ($\psi_i \Leftrightarrow \psi_j$): Consuming one source reward may change the available reward in another sources (e.g. if we need to write all available pencils and pens may become sources of reward, but taking one of them would consume the reward from all sources). Following the previous metaphor, sources seen as wells may be completely or partially sharing their reward pools with other sources. □

These relations can be represented in an interdependence graph G_I . All sources $\psi \in \Psi_W$ are referenced as nodes N in the graph, while edges E represent interdependence relations between sources. Unconnected nodes in this graph represent independent sources (though they may still be subject to changes due to agents' actions or time restrictions). Consequently, the complete world representation W_{SR} may

be defined as follows:

$$W_{SRS} = \{\Psi_W, G_I\} \quad (8)$$

As a toy example, let us imagine a robot buying groceries in a present time supermarket (Fig. 6). All products but one have been collected, and we assume it has to pass through one of the three available cashiers to pay. In this setting we may find a *soft precedence* relationship between the last product and the cashiers' sources. The product may or may not be taken before paying, but it can't be collected after this. The cashier sources have a *codependent consumption* relationship, as they offer the same service once paying on one there remains no reason to go to the others. Finally, for legal reasons it is mandatory to pay before getting out of the market. Therefore, there exists a *strict precedence* between the sources attached to those concepts.

Precedence relationships can be specially useful to shape paths (Fig. 7). While a similar behaviour can be obtained through negative rewards, they can easily provoke inaction, specially if a number of them are stacked in the world representation. This could be solved by resizing the positive reward sources, but it would go against our objective to achieve a modular and additive representation of the world.

3.3 Human–Robot Interaction Using SRS

Humans use abstract concepts as “room” or “flat” on a daily basis, as well as relative quantification of gradable language (e.g. rather, quite, very or dreadfully urgent). Likewise, we

usually use spatial and demonstrative language, often positioning objects in relative references (behind, in front, above, at a certain distance of...). As introduced before, the main motivation behind our spatially centered source definition is dealing with these concepts.

Assuming there exists an ontological knowledge of the environment, it becomes straightforward to encode simple orders or suggestions into sources of reward. “Stay behind me”, “avoid the living room” or “don't stay between me and the tv” can be defined by certain areas (e.g. in \mathbb{R}^2) or volumes (e.g. in \mathbb{R}^3), whose boundaries may be strict or blurry. Such concepts can be represented through a source of reward ψ_{n+1} and, as the SRS Model naturally includes multitask settings, they can be directly introduced into the task world representation.

In these sources, the amount of reward production and R may encode the sense of importance delivered through gradable language or, for example, the propensity to follow someone's instructions. They can be modelled following the designer's criteria and be directly integrated into the world model in an online setting.

$$\Psi'_W = \Psi_W \cup \{\psi_{n+1}\} \quad (9)$$

Most sources of reward generated through human–robot interaction generate either task or agent interaction rewards (depending on the spatial reference). Alternatively, teammates' communication may sprout from a world model reconciliation process. The robot might acquire new infor-

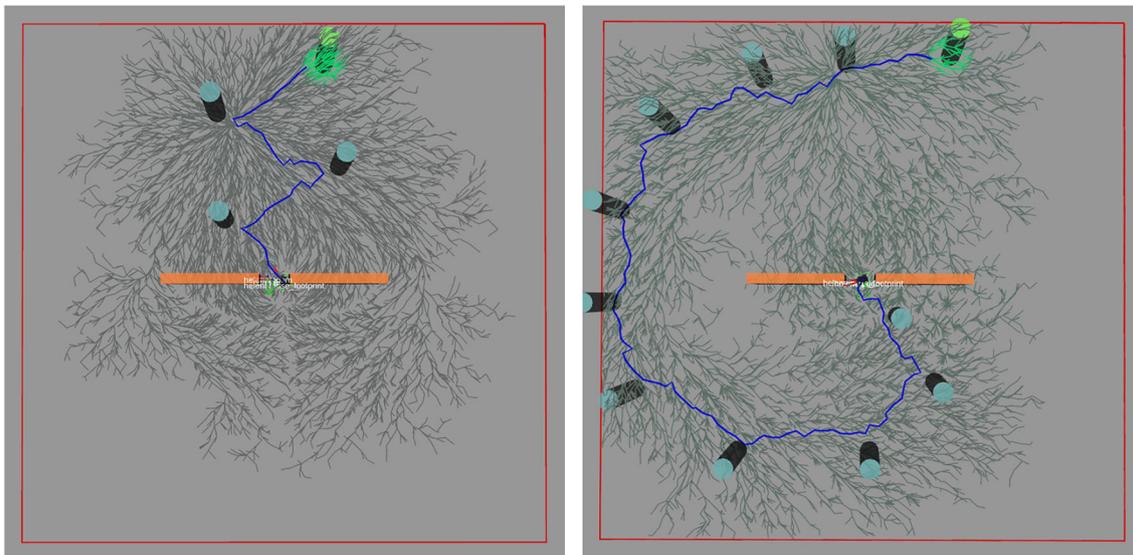


Fig. 7 Path shaping through soft precedence. The plans depicted are generated through expanding an RRT^* in the SRS world. In both examples, all blue sources have a soft precedence relationship with the green one. Due to the soft nature and there being no relationship defined between the blue sources, the path obtained may sometimes skip some

sources or follow a different order to collect them than the one depicted. Note too that the RRT^* used a multi- RRT approach, generating one tree centred on each positive source. (Source: Dalmaso et al. [37]). (Color figure online)

mation from the environment or an agent’s characteristics or state. In such cases, the information should be encoded in the SRS model and added to its respective set.

3.3.1 Shared World Representation

Apart from handling direct instructions, the aim of the SRS Model is to act as a common knowledge representation. If both agents grasp the SRS concept, it may be possible to achieve a shared task representation and, through it, understand each other’s intentions and task contribution. Sharing knowledge and task representations eases maintaining an estimation of others’ knowledge and perception. This encourages the proposal of multi-agent plans taking into account capabilities and preferences of each agent.

All these virtues, however, build over the assumption of the model being intuitive for humans. To test this property, we performed a user study using the model in a previous publication [36]. This study indicated that non-expert humans can understand the main properties of this knowledge representation and use them to design robot behaviours (Fig. 8).

3.4 SRS in Multi-agent Settings

When dealing with multiple agents $\alpha = \{a_0, a_1, \dots, a_{m-1}\}$, the sources of reward should be classified on other criteria. Each source of reward $\psi_k \in \Psi$ may be consumed by multiple agents, from now on identified as the set of targets $\tau_k \subseteq \alpha$. Sources may have from one (private sources) to m targets (global sources).

$$\Psi_W = \{\psi_0^{\tau_0}, \psi_1^{\tau_1} \dots \psi_n^{\tau_n}\} \tag{10}$$

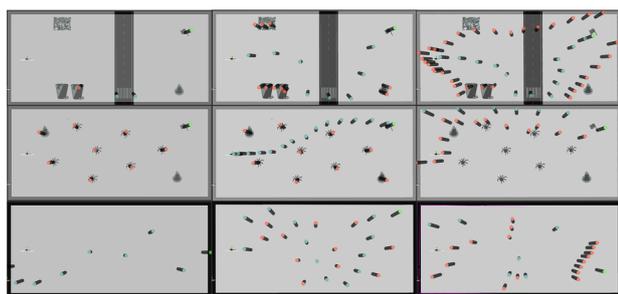


Fig. 8 HRI using SRS. Some designs built by the participants in a virtual user study. In this study, they were empowered to use three instructions: avoid this place (red cylinders), pass through this place (blue cylinders) and go to this place (green cylinders). The scenarios and corresponding indications given to the participants were, following rows from top to bottom. a) Crossroad: Guide the robot to reach the objective (flag) avoiding the objects and crossing the road through the crossroad. b) Spiders: Guide the robot to reach the objective however you feel fitter. c) Free Space: Imagine a trajectory and try to make the robot reproduce it. (Source: Dalmaso et al. [36]). (Color figure online)

Notice this is a representation of shared goals that arises naturally from the source definition. The reward pool of sources having more than one target may be consumed by any of the targeted agents and, consequently, trigger any source interdependence linked to it (Fig. 9).

Taking it one step ahead, one may define the task world representation for each agent $a_i \in \alpha$, each with its reward set:

$$\Psi_{Team} = \{\Psi_{W_0}, \Psi_{W_1} \dots \Psi_{W_{m-1}}\} \tag{11}$$

$$W_{SRS} = \{\Psi_{Team}, G_I\} \tag{12}$$

The sources present in each set would represent the believed knowledge of that agent, i.e. its mental model. Maintaining an estimation of other agents’ knowledge and perception encourages the proposal of plans that take into account capabilities and preferences of each agent.

Imagine how the example depicted in Fig. 9 would change if a_0 believed a_1 didn’t know the light had to be opened so that a_0 could read (i.e. believed a_0 didn’t know about the existence of $\psi_0^{\{a_0, a_1\}}$). Thus, a_0 ’s world representation W_{Team} would include $\Psi_{W_0} = \{\psi_0^{\{a_0, a_1\}}, \psi_1^{\{a_0\}}\}$ and $\Psi_{W_1} = \{\psi_1^{\{a_0\}}\}$. As one solution, a_0 may plan to open the light by himself, as even if he knows a_1 can open the light a_0 doesn’t expect a_1 to do it. Alternatively, a_0 may decide to inform a_1 of the existence of $\psi_0^{\{a_0, a_1\}}$, i.e. engaging in a model reconciliation process.

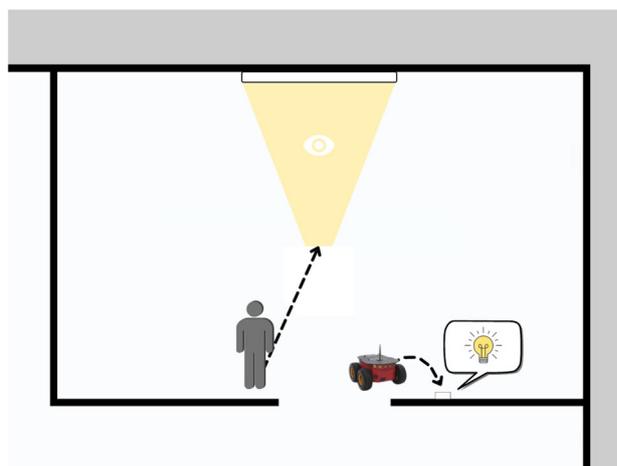


Fig. 9 Multi-agent planning in SRS. Imagine a two agent team whose objective is to read a panel inside the room. Only one of them is capable of reading the panel $\psi_1^{\{a_0\}}$, while both can open the light $\psi_0^{\{a_0, a_1\}}$. There exists a strict precedence between both sources, but a_0 may plan to directly go to the panel as he expects a_1 to take care of the light

4 Multi-agent MCTS

To plan over W_{SRS} we use a decentralized variant of the Monte Carlo Tree Search (MCTS). A thorough explanation of the following formalization can be found in [38].

4.1 Problem Statement

We consider a human–robot team of m agents $\alpha = \{a_0, a_1 \dots a_{m-1}\}$, where each agent $a \in \alpha$ plans its own sequence of future actions \mathbf{x}^a . Each action x_j^a has a required completion time t_j^a . The feasible set of actions \mathbf{x}^a and their associated costs at each step j are a function of the previously taken actions $\mathbf{x}_{0 \rightarrow j}^a = \{x_1^a, x_2^a, \dots, x_{j-1}^a\}$. Thus, for each agent a there is a predefined set \mathcal{X}^a of feasible action sequences \mathbf{x}^a . We will use \mathbf{x} to denote a set of action sequences for all humans and robots $\mathbf{x} := \{\mathbf{x}^{a_0}, \mathbf{x}^{a_1}, \dots, \mathbf{x}^{a_{m-1}}\}$ and \mathcal{X} to denote the set of all feasible \mathbf{x} .

The aim is to maximize a global team objective function $g(\mathbf{x})$ which is a function of the action sequences of all agents. We assume g is deterministic given a known set of action sequences \mathbf{x} , but each agent a may only have a partial or inaccurate perception of $g(\mathbf{x})$, a local objective function $f^a(\mathbf{x})$.

Agents do not know the action sequences selected by the other agents, though $g(\mathbf{x})$ may be modified or $f^a(\mathbf{x})$ updated through teammates communication. As some agents in the team may be humans, the problem must be solved in a decentralised and online setting. Agents may be able to communicate and make shared plans taking into account the objective functions $f^{\{a_0 \dots a_{m-1}\}}$, however they are unable to communicate during planning time to improve coordination. This means they can only interact between the construction of one plan and the next, not being able to change the plan halfway due to new information (some multi-robot approaches [35] include information sharing in the planning process). Therefore, each robot will plan based on the information it has available locally.

4.2 Algorithm Overview

In the presented approach, the objective function g is build through the SRS model. Consequently, the planner must be able to deal with an arbitrary dynamic objective function that may change due to the past action sequences $\mathbf{x}_{1..j-1}$.

The proposed algorithm is designed to run simultaneously and asynchronously on all robots in the team. We will present the algorithm from the perspective of one robot, from now on referred as the agent a_0 . The algorithm cycles between the three phases: (1) incrementally grow each agent search space, a set of feasible action sequences for each agent of the team \mathcal{X}^a , (2) individually compute the probability distribution over each agent possible action sequences and (3) incrementally grow a search tree using MCTS while taking

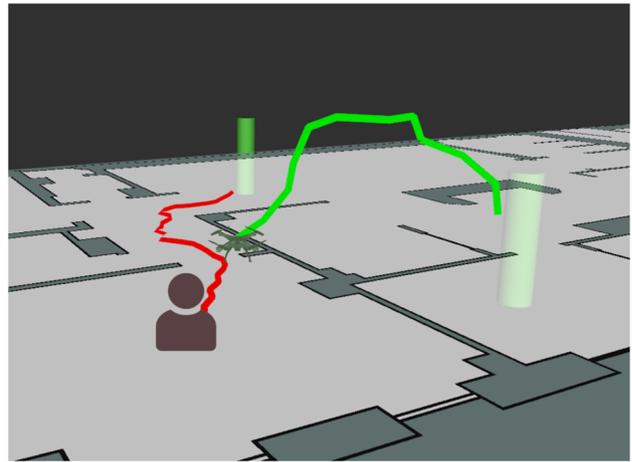


Fig. 10 Heterogeneous search space. In heterogeneous teams each actor may be able to explore different search spaces $\hat{\mathcal{X}}^a$. In the figure a human–drone team is shown, where only the later can explore the space in the vertical dimension and pass above walls. The green cylinders represent sources of positive reward that can be consumed by either agent, indistinctly $(\psi_0^{\{a_0, a_1\}}, \psi_1^{\{a_0, a_1\}})$. The map includes two-meter high walls, but no roof. The given plan is generated by the presented method by taking into account the actors' capabilities. (Color figure online)

into account information about the other agents' objectives and plans.

4.3 Building the Search Space

To generate a feasible restricted set of each agent's possible paths $\hat{\mathcal{X}}^a \subseteq \mathcal{X}^a$ we use a multiple RRT^* . The generated paths can be seen as heterogeneous action sequences, each one leading to a different goal.

Other multi-agent MCTS approaches use a pre-build PRM as the agents' search space (e.g. Best et al. [35]). However, they work with homogeneous teams, where all agents are assumed to have the same mobile and planning capabilities, a premise that does not hold in human–robot settings. Using individual RRTs, in our setting, each agent may have varying capabilities and be subject to different restrictions (Fig. 10).

4.3.1 Agent Action Set

In the present approach, all agent actions are assumed to be movement actions. Each of those are represented by one RRT node i defined by their origin (the node's parent location) and their goal (their own location and its completion time).

Every RRT node i with position \vec{p}_i^a and execution time t_i^a can have an unbounded number of children $ch(i)$, but the number of actions eligible after each agent action is bounded to N_a . Additionally, each action node stores a distribution probability over its children electability on a satisfactory shared plan q_i^a , which is distributed only among the eligible actions.

An achievable reward upper bound is calculated over each individual agent’s action set. Rewards for each agent a action sequences $\mathbf{x}^a \in \hat{\mathcal{X}}^a$ are updated as if they were the only member on the team. Then, all these rewards are back-propagated and each tree node i stores the maximum attainable reward from it. To compute q_i^a , only the N_a children of each node with the highest upper bound are considered, further pruning the search tree. Finally, the selection probability q_i^a is weighted by those bounds.

4.4 Building a Human–Robot Team Plan

The objective of the algorithm is to build a collaborative plan to tackle navigation team tasks over heterogeneous action sets with variable time horizons. To do so, the MCTS planner should contemplate the coexistence of agent action sequences with different temporal length, ensure temporal coherence in the tree expansion and provide a feasible reward propagation mechanism to deal with a dynamic environment.

4.4.1 MCTS State

Each MCTS state s occurs at time t_s , is formed by a list of ongoing agent actions $x_s = \{x_s^{a_0}, x_s^{a_1}, \dots, x_s^{a_{m-1}}\}$ and a list of their remaining time to completion $rt_s = \{rt_s^{a_0}, rt_s^{a_1}, \dots, rt_s^{a_{m-1}}\}$. Note that each agent action x_s^a references one node at that agent’s action set (Fig. 11).

As a state s_k is defined by the ongoing actions x_s , a new MCTS state should be generated after any agent action x_s^a is finished. Also, each MCTS state s_k can have a limited number of children states $ch(s_k)$. Each MCTS state can only have as many successors as the existent possible combinations of finished tasks’ eligible children. Formally, the number of children states is bounded as follows:

$$|ch(s_k)| \leq \prod_a^{\{a_i | rt_{s_k}^{a_i} = 0\}} ch(x_{s_k}^a) \tag{13}$$

4.4.2 MCTS Expansion

The multi-agent plan is expanded from the root node s_0 , which is formed by the root actions of each agent (virtual 0-time initial actions). In each iteration m , one of the expandable states in s_0 is randomly selected through sampling on each finished agent action q_i^a distribution. From the expanded state s_m , a chain of future states is continuously simulated following the same process until a final state is reached. Let us, from now on, refer to this process as rollout. The MCTS tree keeps expanding through continuous rollouts until the predefined time or node limit is reached (detailed in Algorithm 1).

Algorithm 1 $replan(\Psi_W, G_I, \alpha)$

```

1: for each  $a \in \alpha$  do
2:    $\mathcal{X}^a \leftarrow expandRRT(\Psi_W^a, G_I, a)$ 
3:    $\hat{\mathcal{X}}^a, q_i^a \leftarrow pruneActionSet(\Psi_W^a, G_I, N_a)$ 
4:    $\hat{\mathcal{X}} \leftarrow \hat{\mathcal{X}} \cup \hat{\mathcal{X}}^a$ 
5: end for
6: while  $t < t_{max} \parallel n < n_{max}$  do
7:    $S(x_i) \leftarrow rollout(s_0, \hat{\mathcal{X}}, q_i)$ 
8:    $T \leftarrow T \cup S(x_i)$ 
9:   for each  $a \in \alpha$  do
10:     $q_i^a \leftarrow update(x_i^a)$ 
11:   end for
12: end while
13:  $S(x_f) \leftarrow getBestPlan(\Psi_W, G_I, T)$ 
14: return  $S(x_f)$ 

```

Algorithm 2 $rollout(s_k, \hat{\mathcal{X}}, q_i)$

```

1: if  $ch(s_k, \hat{\mathcal{X}}) = \emptyset$  then return  $\{s_k\}$ 
2: end if
3:  $s_{k+1} \leftarrow sample(s_k, q_i)$ 
4: return  $\{s_k\} \cup rollout(s_{k+1}, \hat{\mathcal{X}}, q_i)$ 

```

Finally, note that the q_i^a values related to each agent actions get updated at the end of each rollout. The probability of the actions taking in each rollout is reduced. Such reduction being lower the farther the action is from the current rollout final state.

$$q_i^a(action)' = q_i^a(action) \cdot w \cdot \gamma^n \tag{14}$$

where n is the number of actions between agent a ’s node i and a ’s final action, $w \in (0, 1)$ is the reduction weight and $\gamma \in (0.5, 1)$ is the weight discount factor.

An example of an ongoing collaborative plan expansion of a two member human–robot team may be found in Fig. 11. Additionally, to ensure the viability of early plans we collect potentially rewarding goals from the individual agents’ action sequences and combine them to define end state candidates. This allows for a predefined preliminary expansion to the MCTS tree.

4.5 Objective Function

The collaborative plan objective function $g(\mathbf{x})$ can be seen as the additive combination of all the rewards influencing the team. This includes the rewards related to each agent actions cost A , their perceived influence of the environment E , the tasks at hand T and team interaction I .

$$g(\mathbf{x}) = \sum_a^\alpha (R_A^a + R_E^a) + R_T + R_I \tag{15}$$

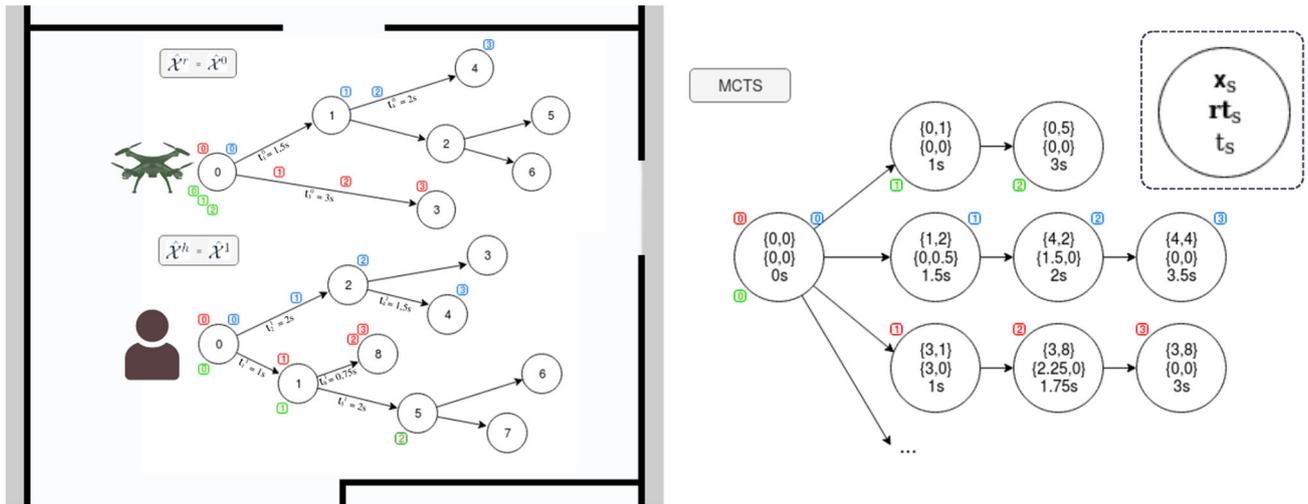


Fig. 11 MCTS tree expansion. Example of a collaborative plan expansion for a two-member team. From left to right: **a** Agent action sets, which constitute the search space of the MCTS. Each node in these sets represent the movement required to reach a determined pose from the final pose of its parent. They store their final pose \vec{p}_i^a , the expected time

to complete the defined action t_i^a and the probability distribution q_i^a . **b** Partial depiction of one possible tree generated by the MCTS. The green, blue and red markers illustrate three possible different plans, whose location on each agent action set can also be seen in figure a). (Color figure online)

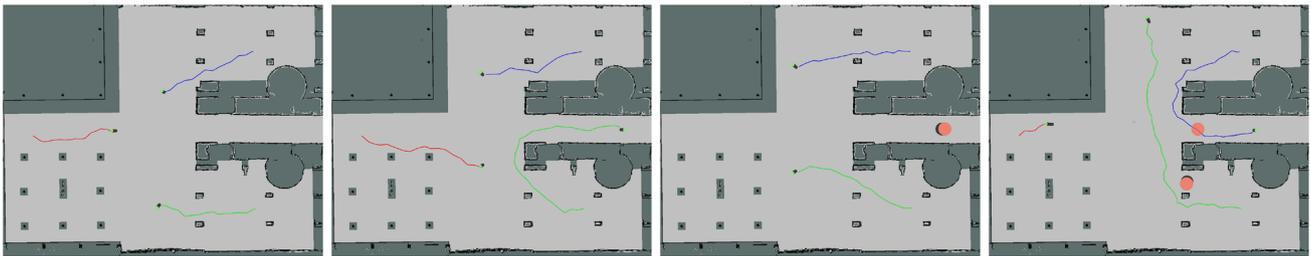


Fig. 12 Human-robot multi-agent MCTS. Planning example for a three-agent team. From left to right: **a** A global plan fulfilling three shared goals. **b** An increase of an individual's effort to maximize global rewards. **c** One of the agents is expected to stay still, as one of the shared

goals is inaccessible due to some hazard (red dot). **d** There are two hazards in the scene, but they only affect one of the agents (green). The planner adapts to fulfil as many shared goals as possible. (Color figure online)

being

$$R_A^a(\mathbf{x}^a) = \sum_{x_j^a \in \mathbf{x}^a} r_A(\Psi_A^a, x_j^a, t_j^a) \quad (16)$$

$$R_E^a(\mathbf{x}^a) = \sum_{s_k}^{S(\mathbf{x}^a)} r_E(\Psi_E^a, x_{s_k}^a, t_{s_k}, \Delta t_{s_k}) \quad (17)$$

$$R_T(\mathbf{x}) = \sum_{s_k}^{S(\mathbf{x})} r_T(\Psi_T^\tau, \mathbf{x}_{s_0 \rightarrow k}, t_{s_k}, \Delta t_{s_k}) \quad (18)$$

$$R_I(\mathbf{x}) = \sum_{s_k}^{S(\mathbf{x})} \sum_a^\alpha r_I(\Psi_{I_a}^\tau, \mathbf{x}_{s_0 \rightarrow k}^{\tau \cup a}, t_{s_k}, \Delta t_{s_k}) \quad (19)$$

where $S(\mathbf{x})$ is the set of MCTS states defined by the action sequences \mathbf{x} . Observe $r_A(\Psi_A^a, x_j^a, t_j^a)$ is the reward collected from the action sources Ψ_A through the execution of action

x_j^a , which takes a time t_j^a . Moreover, $r_E(\Psi_E^a, x^a, t, \Delta t)$ is the reward generated by the environmental sources set Ψ_E^a to target agent a while performing action x^a during a period Δt initiated at time t . Similarly, r_T and r_I are the rewards generated by sources Ψ_T^τ and $\Psi_{I_a}^\tau$ given the action sequences in $\mathbf{x}_{s_0 \rightarrow k}$. Note that $\mathbf{x}_{s_0 \rightarrow k}$ symbolises the action sequences that generate the tree branch connecting the initial state s_0 and the state s_k . Additionally,

$$\Delta t_s = t_s - t_{p(s)} \quad (20)$$

where $p(s)$ is the parent state of s . Consequently,

$$rt_s^a = \begin{cases} rt_{p(s)}^a - \Delta t_s & \text{if } x_s^a = x_{p(s)}^a \\ t_s^a - \Delta t_s & \text{otherwise} \end{cases} \quad (21)$$

The images in Fig. 12 are presented to provide some qualitative example of the model. A human–robot team of three members is given three shared goals and the images show the plan built by one of the robot agents using the multi-agent MCTS model over the SRS representation of four different environments.

5 Collaborative Search

Humans daily engage in searching tasks. Be it keys, writing tools, earring or shoes, we search them following strategies led by our beliefs and use a large spectrum of actions including active perception and object manipulation. Moreover, it’s a situation where we frequently seek collaboration. We have no clear conclusions about the backing motivations for that, but they may include a better expected outcome when adding other’s knowledge and perceptions, better handling of anxiety or trying to avoid other’s unconscious obstruction. During this process, one may update their belief over the object location through the exploration of the environment, the observation of their colleagues’ actions and/or the received information through active communication.

In this work, we consider a constrained search testbed. We assume the environment is fixed. Therefore, there is no occlusion removal through physical interaction. Moreover, the agents’ perception capabilities are constant over time. Specifically, the robot sensors are static in the robot platform reference and the target object is on ground level, considered equally visible on all locations. In short, in this work we consider search as a pure navigation task. Collaborating agents are assumed to be engaged in the task. Moreover, they can move freely, follow independent search policies and lose contact.

5.1 Problem Statement

A human–robot team of agents α explores a known space to locate a set of similar objects O . Each agent may actively search the object, observe their teammates actions and/or communicate with them. The task is assumed as finished when all the objects are found.

In the real-world experiments presented in this work, the team consists of one person and one robot, $\alpha = \{a_0, a_1\}$. Both agents are assumed capable of autonomously navigating through the search area. To construct the robot knowledge of this task, we built an observability graph upon a discretised representation of the search area.

The formal definition of the task detailed in the next subsections was already introduced in [38] and added here for clarity and completeness.

5.1.1 Agent Detection Model

The agents’ object location belief is continuously updated based on the team members’ actions. Hence, we should model the probability of each agent $a \in \alpha$ detecting an object O at a certain location \vec{p} .

$$P(D_a(\vec{p}, \Delta t)|O(\vec{p})) \tag{22}$$

where $\Delta t = t_f - t_0$ is the search time, $D_a(\vec{p}, \Delta t)$ states one object is detected by agent a and $O(\vec{p})$ is the fact that an object is actually at the given location. To build this model, we make a number of assumptions:

Assumption 1 Detection models are independent of their initial time t_{0_i} . In other words, human detection capability does not change over time. So for one agent:

$$P(\overline{D}_a(\vec{p}, \Delta t)|O(\vec{p})) = \prod_i P(\overline{D}_a(\vec{p}, \Delta t_i)|O(\vec{p})) \tag{23}$$

where the overline in \overline{D}_a expresses the complementary statement (i.e. being “undetected by agent a”) and

$$t_{0_{i_0}} = t_0, \quad t_{f_{i_f}} = t_f, \quad t_{f_i} = t_{0_{i+1}} \quad i = i_0, \dots, i_f$$

Assumption 2 Detection models are independent of the agent location. Human detection capability is independent on the perceiving human position, as long as \vec{p} is visible, and on all other participating agents’ position. Whereas, distractions or occlusions in its field of view due to other teammates proximity are not considered. So for each location:

$$P(\overline{D}(\vec{p}, \Delta t)|O(\vec{p})) = \prod_a^\alpha P(\overline{D}_a(\vec{p}, \Delta t)|O(\vec{p})) \tag{24}$$

5.1.2 Object Location Probability

At a given time t , where t_0 is the task beginning and $t_f = t$, the updated object probability on each location given the current accumulated global search is:

$$P(O(\vec{p})|\overline{D}_t) = \frac{P(\overline{D}_t|O(\vec{p})) \cdot P(O(\vec{p}))}{P(\overline{D}_t)} \tag{25}$$

Note that we consider $P(O(\vec{p}))$ as the object probability prior, being $P(O(\vec{p})|\overline{D}_t)$ the object location probability conditioned to the current search (the object not being localised until time t).

To update the global object location belief, we make the following assumption:

Assumption 3 Humans make no false positives while searching or, on another perspective, they filter them automatically. Consequently:

$$P(O(\vec{p})|\bar{D}_t) = \frac{P(\bar{D}_t(\vec{p})|O(\vec{p})) \cdot P(O(\vec{p}))}{1 - \sum \bar{p}_i P(D_t(\vec{p}_i)|O(\vec{p}_i)) \cdot P(O(\vec{p}_i))} \tag{26}$$

The previous formula can be further simplified for the uniform prior case. If we are working on a uniform space discretisation, we may use:

$$P(O(\vec{p})|\bar{D}_t) = \frac{P(\bar{D}_t(\vec{p})|O(\vec{p}))}{N - \sum \bar{p}_i P(\bar{D}_t(\vec{p}_i)|O(\vec{p}_i))} \tag{27}$$

where N is the cardinality of the space discretisation and $P(\bar{D}_t(\vec{p})|O(\vec{p}))$ is iteratively updated from observations using the detection model. One may obtain an efficient belief update calculation using a dynamic programming approach.

5.2 Task Modelling Through SRS

The search task is represented as set of sources Ψ_S with cardinality equal to that of the search space discretisation. Each possible object location L has an equivalent source $\psi_L \in \Psi_S$. Each source ψ_L generates a reward proportional to the probability of detecting the object on the source location along the search space.

$$r(p) \propto P(O(\vec{p}_{\psi_L})|\bar{D}_t) \tag{28}$$

The reward pool of each source is implicitly determined by the detection probability decrease over time (Fig. 13). As the probability of finding the object in one location p nears zero, so will the available reward. The sources defined for the search task present no interdependence between themselves, though all have a *soft precedence* relation to the object localisation event ψ_O .

$$W_S = \{\Psi_S \cup \{\psi_O\}, G_S\} \tag{29}$$

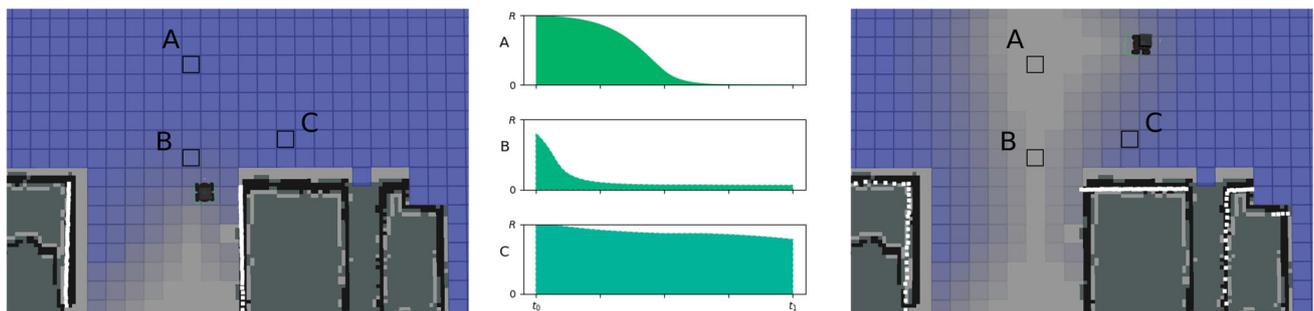


Fig. 13 Reward consumption. Left and right pictures represent the initial and final states of a robot search episode. The central graph depicts the available reward on three selected illustrative lookouts A,B and C, it gets reduced over time as the robot moves up, stops, turns and moves right

H \ R	Human Goal	Robot Goal	R & H Goals	None
Human Goal		L - L		
Robot Goal	P - P			
R & H Goals				L - F (R - H)
None			M - S (H - R)	

Fig. 14 Decision control over planning. This table classifies the arising collaborative planning situations from each agent control over their goals. This paper studies collaboration when agents have partitioned decision sets. Grey cells are undefined behaviours, as the set is undetermined. Conversely, in salmon cells agents have shared decision control over one or both goals during planning. Thus, models tackling the last should include negotiation

$$G_S = \{\Psi_S \rightarrow \psi_O\} \tag{30}$$

Belief over object location probability is continuously updated along the search. As a result, the total reward available is proportional to the believed probability of seeing an object from each given lookout (discretisation block, Fig. 13).

6 Collaboration Roles

Collaboratively planning a search task, as well as the negotiation process that may arise, can have many layers of deepness and insight. As a first approach, we defined a simplified set of decisions that may be taken by either member of a human-robot pair team: the *human goal* and the *robot goal*. Planning occurs many times in the search process as the object location belief of the team gets updated. Consequently, both

goals refer to the next local goal taken in a certain search state. These goals condition the current search plan, but they may or may not be reached if a change in the environment motivates a new planning phase before it happens.

Given the set unavoidable team decisions $D = \{human\ goal, robot\ goal\}$, Fig. 14 classifies the arising collaborative planning situations over each agent control. This paper studies collaboration when agents have partition decision sets ($D_R \cup D_H = D \ \& \ D_R \cap D_H = \emptyset$).

In the presented experiments, the human can communicate with the robot through a set of four different messages. Using the 'My pose' message allows them to indicate to the robot their position, either for re-identification or for conveying information when under occlusion. Secondly, the human can indicate their intention through the 'My goal' message. Similarly, they can also control the robot through the 'Robot goal' message. Finally, the human can also request from the robot a new plan for both the human and the robot sending a 'Replan' message. Doing so, the robot will propose a new plan taking into account all the information conveyed in the previous messages. The replan message may also be used to implicitly accept or reject the plan proposed by the robot by requesting a new plan.

Summarizing, collaboration roles for a human–robot team tackling a search task may vary depending on the decision capability of each agent. Our system allows these relationships to change along the task by empowering the human to do so through a given set of messages. From here on, roles arising with this setting are discussed, while a visual summary of their transitions is depicted in Fig. 15. Note that on the following subsections R stands for *robot* and H for *human*.

6.1 Leader(R)–Follower(H)

A leader-follower relationship occurs naturally if the human sets no restrictions before calling for a Replan. In this state, the robot takes the initiative of proposing a plan for the team, optimising over the task representation. Both human and robot final positions in the plan might change on different executions.

In this mode, the human is still required to indicate their position through 'My pose'. Additionally, the human can always reject the shared plan proposed by the robot by pressing the 'Replan' button again. Even though the robot has complete freedom on the planning phase, both the veto option and the decision of when to request a new plan make us label the relationship as leader-follower, instead of master-slave.

6.2 Peer(H)-to-Peer(R)

A peer-to-peer relationship arises if the human also indicates their intention by using 'My goal' messages. This fixes the

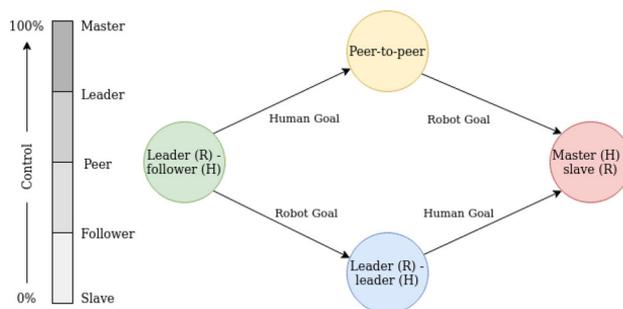


Fig. 15 Role transitions. The left bar summarizes the usual terms to classify collaboration roles as a gradable scale. The right figure depicts the transition graph between team roles in the collaborative search episodes. In the root node the robot decides both goals and transitions occur whenever the human decides to take control over one of the goals determination

human final position and constrains the robot options in the planning phase. The robot has to comply with the restriction and adapts its own planned path to maximise task completion. It must be kept in mind, however, that the robot may still propose different paths leading to the given human goal.

As in the previous case, the human still retains its time control and veto capability through the 'Replan' functionality.

6.3 Master(H)–Slave(R)

We consider a relationship shifting towards master-slave whenever the human decides to indicate both their goal and the robot's. This happens when both 'My goal' and 'Robot goal' messages are used. In this case, the robot still has some flexibility to adapt its trajectory trying to minimise the overlapping of explored areas, but it can be seen as a logic interpretation of the human's orders.

As in all other cases, the human can always reject the shared plan proposed by the robot by pressing the 'Replan' button again, thus having a last level of control over the robot in case they decide to exercise it.

6.4 Leader (H)–Leader (R)

In the defined system, a singular interaction might occur if the human decides to set intention and preference for the robot's actions through 'Robot goal' messages, but leaves 'My goal' undefined. In this case, the robot will propose a team plan complying with the restrictions, but it will also take the initiative in proposing goals for the human.

It may be difficult to imagine a human–human collaboration where this kind of interaction appears. Humans facing a similar situation might simply change roles if possible. Human–robot teams, however, usually present heterogeneous capabilities. It may be easier to find situations where

the human wants to have close control over the robot actions (possibly critical) while they have no objection to leave their own contribution to the robot's optimising process.

Anyhow, this kind of interaction can arise from the given options, so it is taken into account.

7 Smartphone App as HRI Interface

To enable the usage of the SRS model in human-robot communication, we designed a mobile application [39]. It has allowed us to share knowledge and task representations through a user-friendly HRI graphical interface.

Collaborative search tasks can be carried out over arbitrarily large areas, which encourages the appearance of multiple occlusions, high ambient noise and possible resolution problems. Gesture or speech-based communication can be difficult in this setting. On the other hand, using specific gadgets (augmented reality glasses, microphones...) to enable this exchange can be annoying for the user. With all this in mind, we chose to design a mobile application as an HRI interface as we consider it to be the least intrusive yet functional solution.

Based on the previous work from Kohler et al. [30] we have designed an Android application that makes use of both the connectivity and the touch screen of the device. Through this application, the messages delivered by the human are converted into either information or sources of reward that the robot can understand. Likewise, all messages that the robot wants to send to the human are converted into visual information shown on the screen (Fig. 16). Communications are carried out through the deployment of a local Wi-Fi network.

7.1 Messages

On one hand, the interface allows the human to see the map of the current task area. The robot conveys this information to the human phone, as well as their team members' believed location. Lastly, it also shows the last proposed team plan and provides visual feedback on whether the robot is computing a new plan.

On the other hand, the interface allows the human to send messages to the robot. These messages are predefined as: 'Replan', 'My goal', 'Robot goal', 'My Pose' and 'Robot pose'. Knowing the position of both agents is essential both for the robot when calculating the shared plan and for the human when making any decision about which area to explore next. The 'My Pose' message is added for model reconciliation, both as a first identification and as an aid when the location of the human is difficult to track (e.g. under occlusion). The 'Robot Pose' message is used at the beginning of the episodes to approximate robot localisation without moving it.



Fig. 16 Mobile App main screen. *Top*: Screenshot of an ongoing human-robot collaborative search where the human has used every available functionality. 1) Replan button. 2) Input data selection menu. 3) Robot's current position. 4) Human's desired goal for the robot. 5) Path calculated by the robot. 6) Human's current position. 7) Human's intended goal. 8) Path calculated for the human. *Bottom*: Execution using a simplified version of the interface for human-human collaborative search experiments. 1) Input data selection menu. 2) Human's current position. 3) Human's intended goal. 4) Human's partner position. 5) Human's partner intended goal. (Color figure online)

'My goal' and 'Robot goal' functionalities act implicitly as positive reward sources in the SRS world representation. All search sources have a soft precedence relationship to them and they serve to generate rewards that condition the robot's calculation of the routes. Both having more message options and showing the area the robot believes to be explored were discarded for clarity and message bandwidth usage. However, the flexibility of mobile applications allows us to add this or any other information that we deem necessary in future approaches.

7.2 Triggers

The ability of the robot to estimate which areas have been explored can be very dependent on the frequency in which the human indicates his or her position. Also, the robot will need to know the human's position in order to propose a valid plan and, possibly, check if they are indeed following it. This is a common problem in any navigation task with occlusions and this solution requires the user's participation. We took the concept of trigger as an impulse that provokes an action

[32] and used it to design four triggers or reminders with the intention of stimulating the user to indicate their position more frequently.

The first trigger is *'Toast'*, i.e., a small message at the bottom of the screen that tells the user too much time has passed since the last time their position was updated. Namely, that more than a threshold time has elapsed since the last update, which gets reduced if the human doesn't comply. The second, *'Pop-up'*, is a pop-up window that blocks the use of the interface until the user accepts the warning and indicates their position. While the first one is a simple reminder, this one encourages the user to indicate their position before the next warning. The third, *'Multimodal'*, reminds the user to indicate their position through various sensory channels. After the same threshold as in the previous cases, it activates the smartphone's vibrator and blinks the *'My pose'* option on the screen so that the user receives the reminder whether they are looking at the screen or looking at the ground, since they are searching for an object at ground level.

In addition, we developed a fourth trigger using gamification techniques [40]. This consists of adding a counter to each of the available options (*'My pose'*, *'My goal'*, *'Robot goal'*) visible to the user. In this way, the human can know how many times they have used each of the options. The effects of each of the triggers will be analyzed in Sect. 8.4.2.

8 Experiments Field of Study

In this section, we present the experiments that show the performance of the described model.

To our knowledge, as explained in Sect. 3, there are no other methods used in human–robot collaborative navigation that we can compare our system with. While we can provide some references of similar tasks applied to multi-robot settings [24, 35], they are not comparable to our method due to the presence of the human. Similarly, we have come across references to social navigation [22, 41] and other human–robot collaboration approaches [6, 7, 23, 26, 28], but they are task-focused and are not applicable to our experimental setting. Consequently, experiments have been designed to test the suitability and effectiveness of proposed system in the collaborative search HRCN scenario.

This section is structured as follows. First, Sect. 8.1 describes the robot and experimental domain characteristics. Second, Sect. 8.2 describes each experiment design. It is subdivided into Sect. 8.2.1, which describes two interface validation experiments, and Sect. 8.2.2, where the characteristics of the human–robot collaborative search (HRCS) experiment are defined. Following this, Sect. 8.3 lists the demographic data of the participants of each experiment. Similarly, it is divided into Sects. 8.3.1 and 8.3.2, containing participant information related to interface validation

and HRCS experiments, correspondingly. Finally, Sect. 8.4 presents the results of the experiments: interface acceptability is evaluated in Sect. 8.4.1, interface triggers' effect is analysed in Sect. 8.4.2, collaborative search experiments are showcased in Sect. 8.4.3 and the inferred collaborative role taken by the team is discussed in Sect. 8.4.4.

8.1 Robot & Environment Domain

The application has been tested on two different smartphones, the Google Nexus 5 and the Samsung Galaxy S10, both running Android 10 (kernel 4.14. October 2020 compilation). The robot used is based on a Pioneer 2AT platform and the object to be found are three Parcheesi tokens placed on the floor always close to each other at a distance of about 15 cm (Fig. 17). Their position is chosen randomly along the map but more than 10 m away from the initial position of the agents (to ensure a minimum duration of the experiment). As for communications, we have resorted to deploying a local Wi-Fi network using two routers to cover the entire area of the experiment.

All the experiments carried out to validate both the mobile application and our shared planning method were performed at the Barcelona Robot Lab¹ using an area of about 750 m² (8000 ft²). This area is composed of an open zone with multiple occlusions (e.g. walls and columns) and dynamic obstacles such as other passers-by may appear (Fig. 17). The need for the use of our application may become more evident in a larger area, however we opted for a medium-sized one to avoid the volunteers' exhaustion and forfeit. Volunteers are not getting paid and the larger the search area, the longer and harder the experiment will get.

8.2 Experiment Design

Experiment design details are presented in two subsections: *interface validation* and *collaborative search*.

8.2.1 Interface Validation

The first two experiments are used to test the acceptability of the application and to test the effectiveness of the different triggers. This first stage is focused on testing the communication system and is independent of the functioning of the planning algorithm. These experiments have been conducted by human pair teams using a simplified version of the presented application (Fig. 16 - *Bottom*). The objective of the first of these experiments is to study the benefits of using the application, whilst that of the second is to test the effects of the triggers on the participant's behaviour.

¹ <http://www.iri.upc.edu/research/webprojects/pau/datasets/BRL/>.

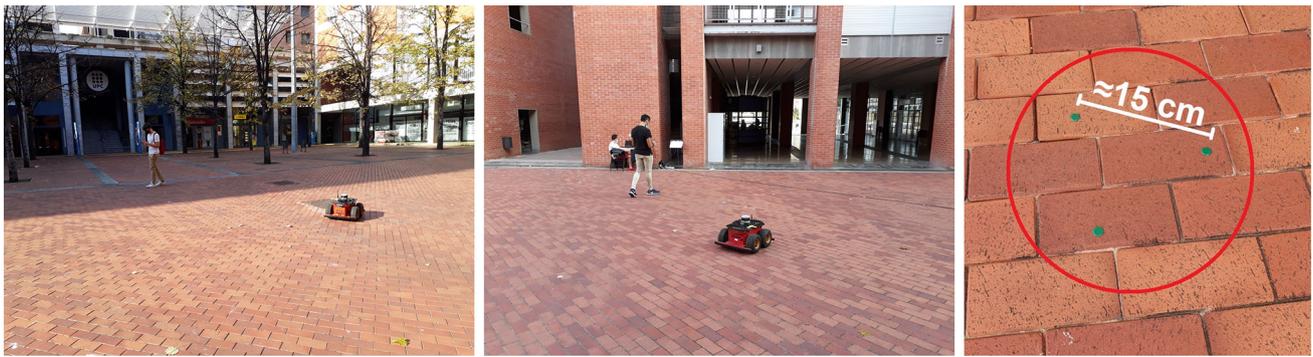


Fig. 17 Experiment examples. *Left & Center:* Human–robot pairs collaboratively searching. *Right:* Group of three green Parcheesi tokens object of the search

Participants in the first interface experiment conducted two different collaborative searches. On the one hand, the first search was done without the app, so they had to communicate with their partner through normal means. Some of the observed means were speaking, shouting, using gestures or establishing an initial plan at the beginning of the search. On the other hand, the second search was performed with our interface running on two mobile phones, one for each member. They were only allowed to communicate through it, which means that they could not use gestures or speak. In this experiment, all volunteers collaborated one by one with the same research assistant, from now on control user, to ensure that the second member of the pair always behaves in the same way. The control user adapted to the information provided by the human as the robot would. Once the volunteers had finished the second search, a questionnaire was given, which we evaluated using ANOVA tests (see results in Sect. 8.4.1).

Participants in the second interface experiment conducted collaborative searches in a setting similar to that of the previous experiment. Each one performed a search with a control user only communicating with him through the interface. They repeated the same collaborative search twice, each time receiving a randomly selected trigger among the four explained in Sect. 7.2. As they were a different population sample, potentially with different biases, preferences and typical behaviours, they also carried out a first search using the application without any trigger in order to obtain a baseline to which compare the results obtained for each trigger. In addition, to avoid the learning effect, a minimum of three days elapsed between the first search, with no triggers, and the other two searches, with triggers. At the end of the experiments, another questionnaire was given. Their answers were again evaluated using ANOVA tests (see the results in Sect. 8.4.2).

8.2.2 Collaborative Search

For the human–robot collaborative search experiments, the full version of our interface as depicted in Fig. 16 - *Top* was used. The '*Robot goal*' function was added in order to specify where the user wants the robot to go, allowing the emergence of the different types of relationships presented in Sect. 6. Also, the '*Replan*' button was incorporated. In order to encourage the user to indicate their pose, the Multimodal trigger is used. Additionally, since the previous two new functionalities increase the amount of time the user spend looking at the interface, we make the human's icon color change from green to red as a visual trigger (equivalent to Toast) to indicate that the robot has lost track of the user.

As in the previous experiments, we tested the model in a two-agent collaborative search environment, although here human participants collaborate with one robot in the search task of three green Parcheesi tokens. To evaluate task progress and performance, the searchable area is virtually discretised and all obstacles in the scene are assumed to block both the view of the robot and the human.

Each experiment episode follows the same operation sequence. The human and the robot begin the search from the same location, on a side-by-side initial pose. First, the user indicates their position allowing the robot to identify them as its search partner. The experiment begins once the user presses the Replan button, which triggers the calculation of a team plan by the robot. Once the plan is available, it's shown on the user interface. At this point, if the user dislikes the proposed plan, they can reject it by requesting a new plan through the Replan button. Otherwise, the plan will be considered as accepted and the robot will start moving.

At any given moment, even previously to the first replan, the user may introduce restrictions to the plan specifying their goal and/or the robot's. As introduced before, this changes

the nature of the interaction roles. Plans may also be rejected halfway by the human or recalculated by the robot if any obstacle to their completion is found. This process is repeated until the Parcheesi tokens are found. All listed functionalities can be observed in the explanatory video.²

8.3 Participants

A total of 71 different volunteers participated performing up to 153 experiments. Participants on each experiment are a different population, so their demographic data is reported separately in their corresponding subsection (see 8.3.1 and 8.3.2).

All the experiments reported in this document have been performed under the approval of the ethics committee of the Universitat Politècnica de Catalunya (UPC)³ in accordance with all the relevant guidelines and regulations. All volunteers who have participated in these experiments are of legal age and in full use of their mental faculties. All of them have signed an informed consent form after having received all the relevant information regarding the experiment and before the beginning of the experiment. Additionally, they have accepted that all the information collected during the experiments (messages exchanged through the application, sensor readings, answers to questionnaires, photographs, and videos) will be treated anonymously for academic purposes.

8.3.1 Interface Validation

In the first interface experiment there participated 20 volunteers, 7 women and 13 men, between 18 and 26 years old (mean: 21.1, std: 3.46) and with an average education level of B.Sc. (as most common ongoing or finished studies). Their self-evaluated knowledge in robotics was 2.90 (std: 1.25) on a scale of 1 (None) to 7 (Expert) and none of them used the interface before. In total, they performed in 40 interface validation experiments.

Likewise, 30 different volunteers, 5 women and 25 men, between 20 and 41 years old (mean: 27.8, std: 5.21) and on average with M.Sc. (as most common ongoing or finished studies) participated in the second interface experiment. Their self-evaluated knowledge in robotics was 4.90 (std: 1.26) on a scale of 1 (None) to 7 (Expert) and none of them used the interface in the previous round of experiments. In total, 90 trigger effect experiments were performed.

8.3.2 Search Experiments

A total of 19 volunteers, 4 women and 15 men, participated in the experiment. They were between 18 and 40 years old

(mean: 21.95 std: 5.39), their education level was on average B.Sc. and their average and their self-evaluated knowledge in robotics was 2.37 (std: 1.21) on a scale of 1 (None) to 5 (Expert). No one could practice using the setting or the mobile app, neither had they any previous experience with it. Each of them participated in one or two episodes. In total, 21 human–robot collaborative search experiments were performed.

8.4 Results

Results from the previously described experiments are discussed in this section. They are presented in four main blocks: *app's acceptability study*, *trigger's effects*, *collaborative search* and *types of collaboration*.

8.4.1 App's Acceptability Study

With the first experiment, our objective was to confirm two hypotheses. First, the user does not reject its use, accepting that it has advantages and disadvantages with respect to other communication channels. Secondly, using our interface as the only method of communication does not reduce the performance of the task or, in other words, that the user is able to adapt to it to perform the task with the same effectiveness or even better.

To test the first two hypotheses, it is necessary to conduct a user study and take not only subjective measures of user experience perception, but also objective measures of task duration and success rate. Figure 18 shows the results of the post-experiment questionnaire answered by the users.

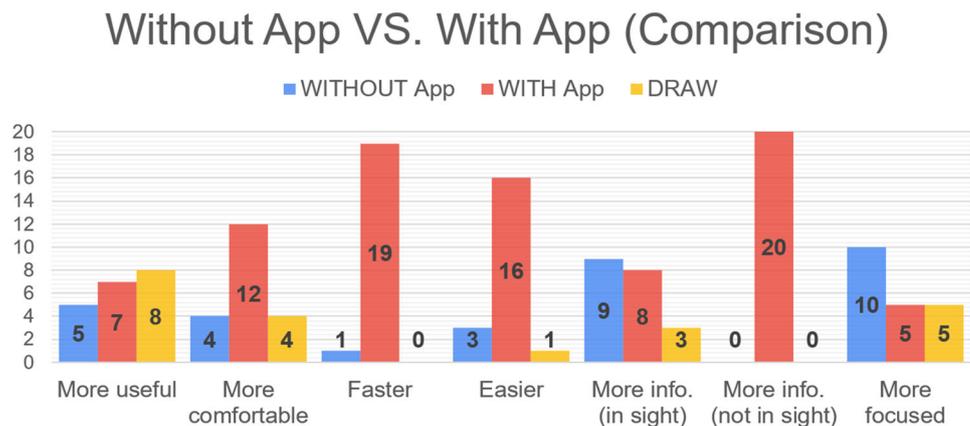
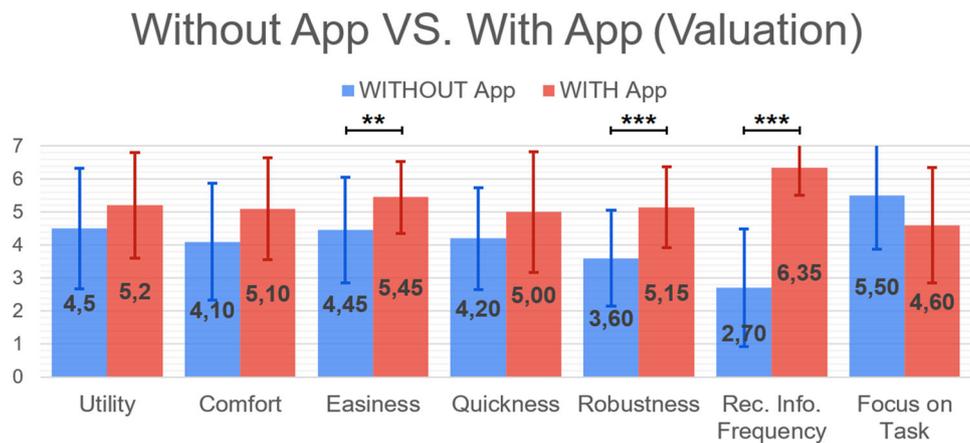
The answers from the users suggest that our interface does not worsen the speed of information exchange, but does make it easier to perform such exchange. Likewise, it seems our interface allows the exchange of more information than the typical methods used in human communication for this type of task. The drawback is that the use of our application shows a concentration reduction tendency, though it is not statistically significant. There is unanimity on the application offering more information when the partner is not in sight. Moreover, a large majority consider that our interface is an easier way of communication. We consider all this confirms our first hypothesis: the users have no inconvenience in using our app and accept it due to its advantages, although they realize it makes it harder to focus on the task at hand.

Looking at objective data, the mean search time of the experiments without the app was 129.1 s (std: 86.0 s), whilst the mean duration of the ones with the app was 124.8 s (std: 79.5 s). Although the difference is not statistically significant, note that the variability is lower while using the app and its usage takes time from the users. In addition, 95% of the searches performed without the app resulted in successfully finding the object while it was a 100% reached while using

² Collaborative search video: <https://youtu.be/k4j51308Zz4>.

³ Ethics committee URL: <https://comite-etica.upc.edu/en>.

Fig. 18 Acceptability user study. *Top*: Valuation of most relevant factors from 1 (low) to 7 (high) regarding their perception of the use of the interface. Statistical significance marked with *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. *Bottom*: Result of making the users to chose among using or not the app (or draw) for the same factors. The maximum score is 20 since there were 20 volunteers in this round of experiments



the app. We consider this to confirm our second hypothesis: the use of our interface does not hinder the performance of the task.

Finally, let us take a look at the exchange of 'My pose' messages which, as mentioned above, are essential for the robot to be able to use its shared planning algorithm. Table 1 shows the clear difference between our control user (the research assistant acting as theatrical robot) and the volunteers (users without previous training and typically with low knowledge about the robot needs). Both if we look at the evolution over the experiments or just at their average, the control user sends about a 125% more messages than the normal user.

A problem may arise if the frequency with which the position is updated is not high enough to be able to interpolate

the path that the user is following. This minimum threshold is relative to the task. In our case, given the volunteers tend to move at a speed which ranges from 0.5 (between columns or near to edges) to 1.0 m/s (in open space), we establish 10 s as the maximum threshold for the time between messages (TBM) of 'My pose' type. Table 1 also shows the percentage of 'My pose' messages that both users send without exceeding this threshold counting since the previous message was sent. Even the control user does not manage to stay above 95% since he is human and sometimes gets distracted. However, the real problem is that the standard user has a 13-23% worse performance making that between 18% and 29% of their 'My pose' messages do not comply with this threshold being this problematic for the robot.

Table 1 Evolution of 'My pose' msgs. rate [msg/min] and time between messages (TBM) < 10 s [%]

User		Experiment evolution			Complete experiment
		60 s	90 s	120 s	
Control	Rate [msg/min] (std. dev.)	17.91 (4.32)	16.61 (3.53)	14.92 (3.57)	15.59 (4.86)
	TBM < 10 s [%]	94.79	94.36	94.25	90.95
Volunteer	Rate [msg/min] (std. dev.)	8.06 (4.54)	7.90 (4.21)	6.08 (3.12)	6.38 (4.21)
	TBM < 10 s [%]	81.82	81.76	70.59	72.94

8.4.2 Trigger's Effects

The second experiment has the objective of confirming one hypothesis: we can encourage the user to increase the frequency with which they indicate their position by using the triggers previously discussed in Sect. 7.2.

As the 30 volunteers in this second round of experiments are a different population sample, we perform a first experiment with all of them and without any trigger to know which one is the threshold which generates similar percentages obtaining that we need to reduce that threshold from 10 s to 8.32 s (76.78% of the messages were sent on average within 10 s with the previous population sample and 76.82% within 8.32 s with this second population sample). Having done this, each volunteer performed two more collaborative searches using in each one a randomly selected trigger.

The effectiveness of each trigger in encouraging the user to indicate their position is shown in Fig. 19. The first trigger, Toast, decreases the frequency in which the users indicate their position. This possibly happens because the users expect to be reminded but this trigger is too subtle for them to notice. Conversely, Pop-up and Multimodal are both quite effective increasing the percentage of messages along the complete experiment, increasing messages with $TBM < 8.32$ s in a 4.89% and 6.92%, respectively. Finally, Counters is promising since with very little computational cost it can give almost the same results with an increment of 5.09% for the complete experiment, although with a lower performance than Pop-up or Multimodal in the early stages of the experiments. This seems to indicate that the users tend to get tired if they are interrupted too often, turning gamification techniques the best option for long-term tasks. Although without reaching the control user level, this data suggests the veracity of our third hypothesis.

As we did to test the degree of acceptance of our interface, in this second round of experiments we also asked users to

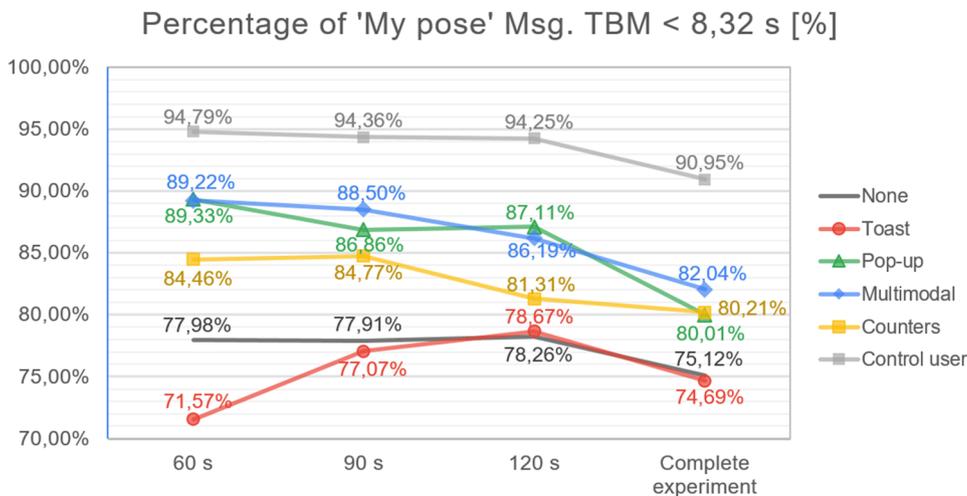
fill out a post-experiment questionnaire to obtain subjective data on their perception of how the triggers work (Fig. 20).

Neither the perceived difference in information sent or received are statistically significant, the second being expected as it is the information sent by the assistant. Concerning their perceived frequency of the trigger reminders, there are statistically significant differences between Pop-up and Multimodal (M) compared to Counters. This means that the users consider that both Pop-up and Multimodal are much more present being Counters more subtle. There is also a significant difference between the perceived frequency in what they do update their position after being reminded. Data suggests users eventually tire of Pop-up and stop reacting to it (as can be seen in the final stages of the experiment in Fig. 19). Similarly, Pop-up is rated as the most annoying trigger, followed by Multimodal. Expectedly, using Counters allows the users to concentrate on the task significantly better than using Pop-up.

Following the same methodology, and anticipating considerable variances as in the previous valuation, we also asked the users to choose between the triggers they used (Fig. 21). The general feeling is that they consider Toast as the least effective trigger and Multimodal as the most effective one, confirming the results obtained in the previous valuation. Regarding the intrusiveness, i.e. how much the trigger interrupts you, they consider that Multimodal and Pop-up tie in this aspect but, if we look at the degree of annoyance, they consider Pop-up as notably more annoying which means that they accept the intrusiveness of Multimodal but refuse the one of Pop-up. Finally, they consider Counters and Multimodal as the best options to keep their attention on the task being *draw* the more common answer when both triggers are compared in this aspect.

To sum up, Multimodal is the most balanced trigger being effective and not too much distracting for tasks which last for about two or three minutes like in our case. However,

Fig. 19 Evolution of each trigger's $TBM < 8.32$ s. Percentage of 'My pose' messages within the 8.32 s TBM threshold using each trigger. Baseline case with no triggers and control user case also shown for comparison



Triggers effects

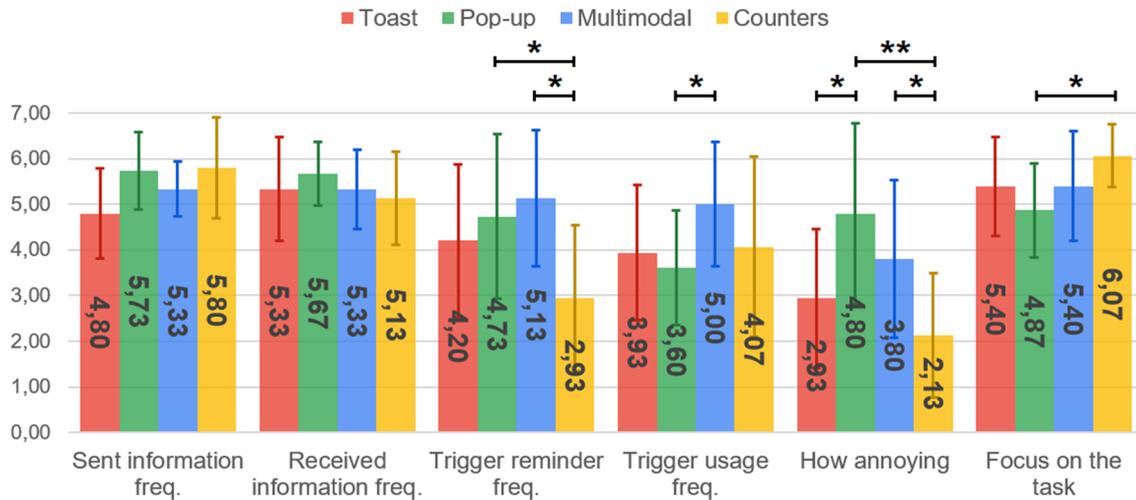


Fig. 20 Main aspects evaluated in the triggers user study. User's valuation from 1 (low) to 7 (high) of the main aspects related to the effects of using each of the designed triggers. Statistical significance marked with *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$

More effective?		Trigger				
	T	P	M	C	Avg.	
Trigger	T	1.5 (3)	0 (0)	1 (0)	0.83	
P	3.5 (3)	2 (0)	3 (0)	2.83		
M	5 (0)	3 (0)	3 (0)	3.67		
C	4 (0)	2 (0)	2 (0)	2.67		

More intrusive?		Trigger				
	T	P	M	C	Avg.	
Trigger	T	0 (0)	1 (2)	4 (0)	1.67	
P	5 (0)	2 (0)	5 (0)	4.00		
M	4 (2)	3 (0)	5 (0)	4.00		
C	1 (0)	0 (0)	0 (0)	0.33		

More annoying?		Trigger				
	T	P	M	C	Avg.	
Trigger	T	0 (0)	2 (2)	3 (0)	1.67	
P	5 (0)	5 (0)	5 (0)	5.00		
M	3 (2)	0 (0)	4 (2)	2.33		
C	2 (0)	0 (0)	1 (2)	1.00		

Easier to focus?		Trigger				
	T	P	M	C	Avg.	
Trigger	T	5 (0)	2 (0)	1.5 (1)	2.83	
P	0 (0)	1 (0)	0.5 (1)	0.50		
M	3 (0)	4 (0)	3 (2)	3.33		
C	3.5 (1)	4.5 (1)	2 (2)	3.33		

Fig. 21 Comparison of triggers. T = Toast, P = Pop-up, M = Multimodal, C = Counters. Scoring as follows: +1 each time a user selects one specific trigger for each factor. +0.5 for both triggers if they select draw. Amount of draws near to the score in parentheses. Each row represents

experiments with longer duration, we have observed that Counters, i.e. a gamification-based trigger, provokes less fatigue on the user as shown in Fig. 19.

8.4.3 Collaborative Search

The human–robot search experiment has the objective of confirming three hypotheses. First, the fact that the SRS model is flexible and expressive enough to represent a complex human–robot collaborative navigation task. Second, given the first hypothesis holds true, the proposed planner is capable of generating feasible and meaningful HRC plans along the SRS world representation. Finally, the human participants in the experiments will have a positive perception of the robot, its capabilities and their safety.

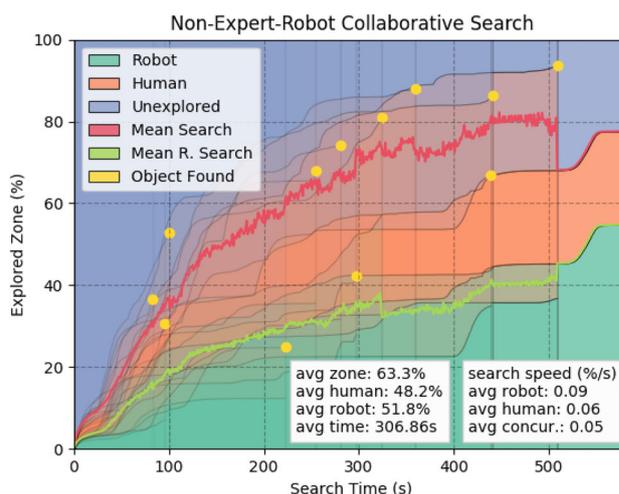
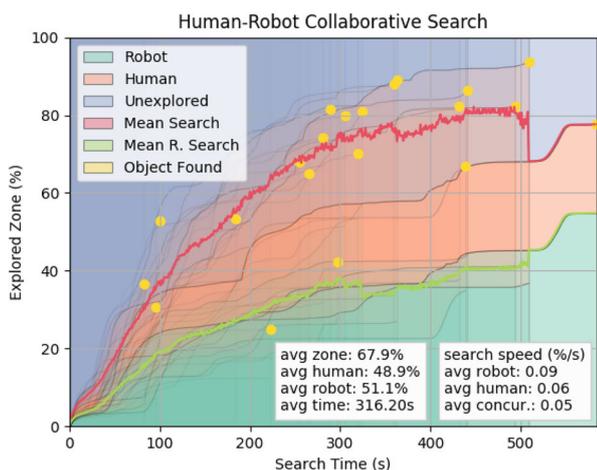
The 21 search episodes from the experiments are summarized in Fig. 22 - Top. Here, each episode is depicted by two stripes: a green stripe that represents the area explored by the robot and an orange stripe showing the perceived contribution of the human. All 21 episodes are stacked in an overlapping fashion and the green and red thick lines represent the mean robot and global exploration progress respectively. In general, the graph represents the task progress over time and each episode finishes when the object is found (see the yellow dots).

resents the comparison between the trigger at the beginning of the row and each other trigger. The last column is the average of the scoring versus every other trigger. There were 30 volunteers and each of them used 2 triggers, so the maximum possible score is 5.0

The search task was completed on all episodes (episodes with technical difficulties were discarded). Episodes presented a mean time to completion of 316.2s and, on average, the team had explored a 67.9% of the map when the object was found. In the robot's perception, both the human and the robot explored more or less equivalent areas. The human contribution, however, might be underestimated due to the asynchronous update of their location in occlusion situations. Moreover, the observable area from the initial location is usually counted as explored by the robot, as it may take the first seconds of the episode to identify the human as the current teammate. Consequently, it can be observed that the robot contribution increases quicker at the beginning of the search.

Given all of the above, we consider as confirmed both the first and second hypotheses, the SRS model is capable of representing a human–robot collaborative navigation task and the multi-agent MCTS planner can find feasible and meaningful solutions over it. Data in Fig. 22 - Bottom also suggests that the third hypothesis holds true.

The collaborative search task was taken both by professionals working in the field of robotics and other people. A comparison between the episodes of both collectives can be observed in Fig. 23. The episodes with the collaboration of roboticists seem to present less variability, but also fewer early-time discoveries.



Robot		mean	std
Likeability	Dislike (1) - Like (5)	4.2	0.9
	Unfriendly (1) - Friendly (5)	4.1	1.2
	Unpleasant (1) - Pleasant (5)	3.8	1.1
Robot Plans	Awful (1) - Nice (5)	4.1	0.7
	Incompetent (1) - Competent (5)	4.0	0.9
	Ignorant (1) - Knowledgeable (5)	3.8	0.9
	Irresponsible (1) - Responsible (5)	4.1	0.8
Perceived Safety	Unintelligent (1) - Intelligent (5)	3.9	1.0
	Foolish (1) - Sensible (5)	3.7	0.9
	Anxious (1) - Relaxed (5)	4.1	1.3
	Agitated (1) - Calm (5)	4.1	1.3
	Quiescent (1) - Surprised (5)	3.3	1.4

Fig. 22 Multi-agent human–robot collaborative search. *Top*: Experimental data from human–robot collaborative search episodes. *Bottom*: Participants’ feedback concerning the robot after the collaborative search experiments. Some of the questions are inspired from the *God-speed questionnaire* [42] (Source: Dalmasso et al. [38])

8.4.4 Collaboration Roles

The messages from the data of the third experiment were extracted to study the type of collaboration that the participants shared with the robot. Following the criteria established in Sect. 6, the collaboration roles have been assigned along the episodes subject to the messages used by the human. A graphical depiction of these can be observed in Fig. 24. The most meaningful area of the graph could be the 100 – 200 s interval. Here, the 85, 71% of the episodes are still ongoing and the role changes from the initial phase seem finished. Both leader-follower and master-slave strategies are observed to be predominant, however neither the peer-to-peer nor the leader-leader strategies have a non-negligible percentage.

As done in the previous section, a comparison between non-expert and expert participants is observable in Fig. 25. It is interesting to note that roboticists actively avoided the leader-leader strategy, possibly because it is an uncommon

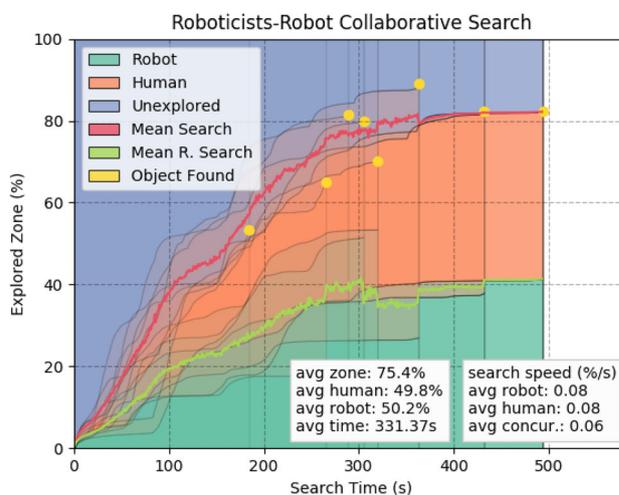


Fig. 23 Non-experts vs roboticists. Experimental data from the human–robot collaborative search experiments split between expert and non-expert participants

human–robot relationship in the field. The strategy was, however, followed by a significant part of the non-expert participants. It may have been wrongly overlooked in the literature and perhaps further research in this direction is needed.

9 Discussion

The SRS model proposes to generalise the world view of the robots (and their perception of the humans’ view) as a set of reward sources. Using rewards to represent the world is not an innovative approach in itself. For instance, sampling methods have been using costs for a long time (arguably negative rewards) and many learning approaches use rewards as their main world feedback for their actions. What might be innovative, however, is the modelling of the sources of such rewards.

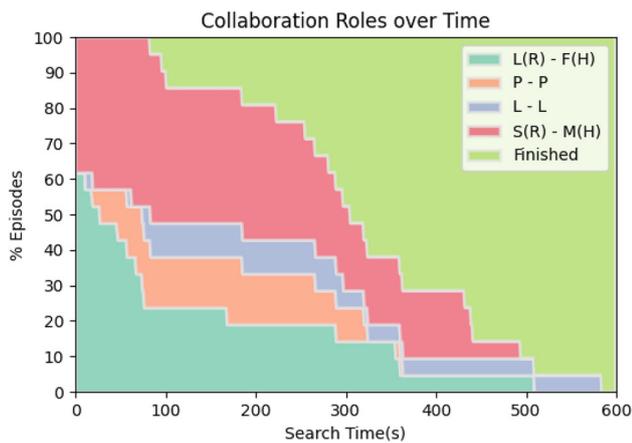


Fig. 24 Collaboration roles. Collaboration roles over time in human–robot search experiments

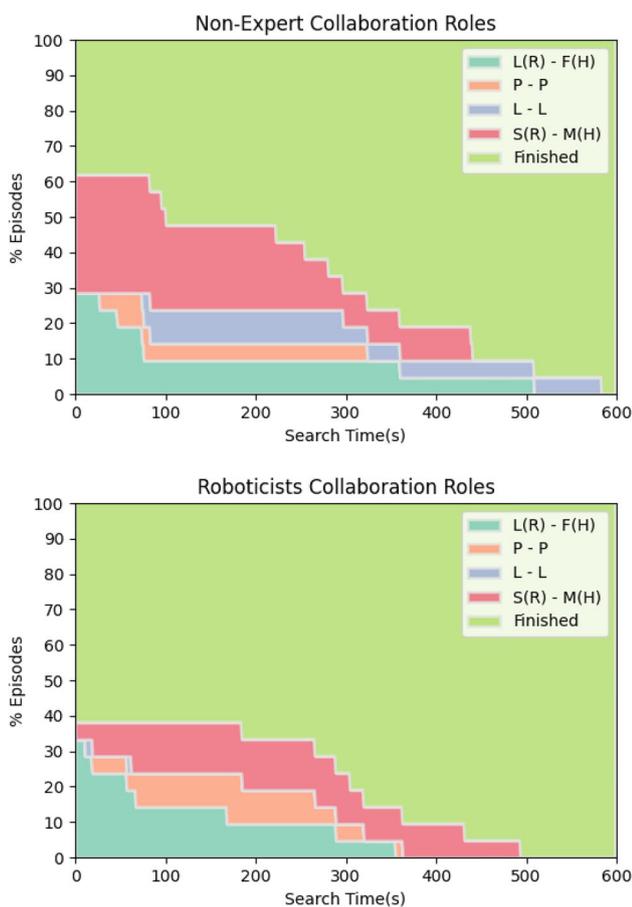


Fig. 25 Non-experts versus roboticists. Collaboration roles from the human–robot collaborative search experiments split between expert and non-expert participants

This paradigm shift draws the world as a modular construction of entities which may or may not be independent of each other and can represent world and task progression dynamics. Both multitasking and interdependency enable planners to consider task allocation and task planning at the same level as path planning. Grounding these usually high-level plan-

ning decisions to the path planning level makes it possible to make decisions over more accurate time and cost approximations in navigation tasks.

A concern that arises over the usage of this model is the scalability. Though it naturally integrates changes in the environment through the addition of new sources of reward, the cardinality of these entities might entail computation problems. This issue might be solved through planning horizon limitations or the generalisation of source sets in individual sources. Further discussion is needed in the world modelling process under SRS. Reward overlapping, a core feature that enables the representation of multifinal actions in multitasking settings, may produce undesired artefacts. Should a cabinet or a table be considered sets of sources representing their different parts, zones with overlapping rewards may overrepresent the designed intent. This discretisation design problem is inherited from previous approaches, like the SFM, and might be problematic when trying to represent human body parts or fuzzy gradable hazards (i.e. dangerous temperatures, gas leakages or spread fire).

The planning approach presented in this work is capable of generating multi-agent plans for heterogeneous human–robot teams over the SRS world representation. This affirmation has been proven through real-world experimentation over the proposed collaborative search testbed. We decided to use RRTs to build the agent action sets for two main reasons: (1) we focus on changing dynamic environments and decided the action sets should be rebuilt (or reconstructed) in each planning phase and (2) constructing independent RRTs allowed us to represent the heterogeneous nature of the human–robot teams and individually consider sources' targets and their effects on the trees' shape. It should be noted, however, that RRTs might not converge into optimal solutions under positive rewards influence and the resulting trees won't offer a complete representation of the agent's options, even at a qualitative level (Fig. 4). In worse cases, global rewards attainable by two or more agents can provoke the generation of overlapping agent action sets where all available team plans involve either more than one agent trying to collect the same reward or some team member's inaction. Perhaps using agent-linked PRMs and dynamically reconstructing them could solve some of these issues, though the existence of cycles in the action set would make the model more complex, with special concern on the UCT computation and the horizon limitation of the MCTS rollouts.

Lastly, using a mobile app for communication in a search task has one main drawback: the participant's attention gets divided between the task and the device. Other technical approaches like using augmented reality headsets might be able to ease handling both simultaneously. Conversely, natural language or gestural interaction could prove to be difficult or impossible when separated in a noisy open space, especially if under occlusion. Consequently, we decided the best

option available to test a possible day-to-day real-world scenario was to use a smartphone interface.

10 Conclusion

The overall goal of this work is to build a system capable of handling generic collaborative navigation tasks. To test this system, human–robot pairs were asked to tackle a collaborative search. The collected data from such experiments proves the SRS model is capable of representing a human–robot collaborative navigation task. Likewise, the multi-agent MCTS planner can find feasible and meaningful solutions over such representation. Additionally, participants in these experiments show a positive perception of the robot, its capabilities and their safety.

The interface used for human–robot communication does not hinder the performance of the team and it is possible to encourage its desired usage using triggers. Concretely, human users have no inconvenience in using our app, accepting it due to its advantages though they realise it makes it harder to focus on the task at hand.

In conclusion, we present a system capable of representing and handling collaborative navigation tasks in human–robot team settings. Such system has been tested online in real-world experiments and the collaboration roles arising from the collaboration are identified and discussed. To our knowledge, there are no similar works in the present literature. We find it promising to test this system in other collaborative navigation tasks (e.g. search and retrieval, side-by-side navigation...) or other team settings (e.g. UAV-human, teams including more than one human and/or more than one robot...). Moreover, we hope the room for improvement in its building blocks to be a stepping stone for future research.

Funding Open Access funding provided thanks to the CRUE-CSIC agreement with Springer Nature. This work was supported under the Spanish State Research Agency through the Maria de Maeztu Seal of Excellence to IRI (MDM-2016-0656) and ROCOTRANSP project (PID2019-106702RB-C21/AEI/10.13039/501100011033), the European research grant TERRINet (H2020-INFRAIA-2017-1-730994) and by JST Moonshot R & D Grant Number JPMJMS2011-85.

Data Availability Available data of this project is in this manuscript. Authors may be addressed for further information.

Declarations

Competing interests There are none potential conflicts of interest that could bias the evaluation or results of our research.

Informed Consent Informed consent was obtained from all individual participants included in the study. The authors affirm that human research participants provided informed consent for publication of the images in Fig. 17a, b. As explained in Sect. 8.3, all experiments were under the approval of the ethics committee of the Universitat Politèc-

nica de Catalunya (UPC) and in accordance with all its guidelines and regulations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Liu R, Natarajan M, Gombolay MC (2021) Coordinating human–robot teams with dynamic and stochastic task proficiencies. *ACM Trans Hum–Robot Interact (THRI)* 11(1):1–42. <https://doi.org/10.1145/3477391>
2. Fan J, Zheng P, Li S (2022) Vision-based holistic scene understanding towards proactive human–robot collaboration. *Robot Comput Integr Manuf* 75:102304. <https://doi.org/10.1016/j.rcim.2021.102304>
3. Costa GDM, Petry MR, Moreira AP (2022) Augmented reality for human–robot collaboration and cooperation in industrial applications: a systematic literature review. *Sensors* 22(7):2725. <https://doi.org/10.3390/s22072725>
4. Chakraborti T, Kulkarni A, Sreedharan S, Smith DE, Kambhampati S (2019) Explicability? legibility? predictability? transparency? privacy? security? the emerging landscape of interpretable agent behavior. In: *Proceedings of the international conference on automated planning and scheduling*, vol. 29, pp 86–96 (2019). <https://doi.org/10.1609/icaps.v29i1.3463>
5. Peternel L, Tsagarakis N, Caldwell D, Ajoudani A (2018) Robot adaptation to human physical fatigue in human–robot co-manipulation. *Auton Robot* 42(5):1011–1021. <https://doi.org/10.1007/s10514-017-9678-1>
6. Liu X, Li G, Loiano G (2022) Safety-aware human-robot collaborative transportation and manipulation with multiple MAVs. *arXiv preprint arXiv:2210.05894*
7. Fontaine MC, Nikolaidis S (2022) Evaluating human–robot interaction algorithms in shared autonomy via quality diversity scenario generation. *ACM Trans Hum–Robot Interact (THRI)* 11(3):1–30. <https://doi.org/10.1145/3476412>
8. Tuli TB, Kohl L, Chala SA, Manns M, Ansari F (2021) Knowledge-based digital twin for predicting interactions in human–robot collaboration. In: *2021 26th IEEE international conference on emerging technologies and factory automation (ETF A)*, pp 1–8. <https://doi.org/10.1109/ETF A45728.2021.9613342>
9. Rudenko A, Palmieri L, Herman M, Kitani KM, Gavrila DM, Arras KO (2020) Human motion trajectory prediction: a survey. *Int J Robot Res* 39(8):895–935. <https://doi.org/10.1177/0278364920917446>
10. Kemény Z, Váncza J, Wang L, Wang XV (2021) In: Wang L, Wang XV, Váncza J, Kemény Z (eds.) *Human–robot collaboration in manufacturing: a multi-agent view*, pp 3–41. Springer, Cham. https://doi.org/10.1007/978-3-030-69178-3_1
11. Lee M-L, Behdad S, Liang X, Zheng M (2022) Task allocation and planning for product disassembly with human–robot collaboration.

- Robot Comput Integr Manuf 76:102306. <https://doi.org/10.1016/j.rcim.2021.102306>
12. Nunes E, Manner M, Mitiche H, Gini M (2017) A taxonomy for task allocation problems with temporal and ordering constraints. *Robot Auton Syst* 90:55–70. <https://doi.org/10.1016/j.robot.2016.10.008>
 13. Johnson M, Bradshaw JM, Feltovich PJ, Jonker CM, Van Riemsdijk MB, Sierhuis M (2014) Coactive design: designing support for interdependence in joint activity. *J Hum Robot Interact* 3(1):43–69. <https://doi.org/10.5898/JHRI.3.1.Johnson>
 14. Elliot AJ, Thrash TM (2002) Approach-avoidance motivation in personality: approach and avoidance temperaments and goals. *J Pers Soc Psychol* 82(5):804. <https://doi.org/10.1037/0022-3514.82.5.804>
 15. Helbing D, Molnar P (1995) Social force model for pedestrian dynamics. *Phys Rev E* 51(5):4282. <https://doi.org/10.1103/PhysRevE.51.4282>
 16. Zanlungo F, Ikeda T, Kanda T (2011) Social force model with explicit collision prediction. *EPL (Europhys Lett)* 93(6):68005. <https://doi.org/10.1209/0295-5075/93/68005>
 17. Wykowska A, Chellali R, Al-Amin M, Müller HJ et al (2014) Implications of robot actions for human perception. How do we represent actions of the observed robots? *Int J Soc Robot* 6(3):357–366. <https://doi.org/10.1007/s12369-014-0239-x>
 18. Lin Z, Zhao L, Yang D, Qin T, Liu T-Y, Yang G (2019) Distributional reward decomposition for reinforcement learning. In: *Advances in neural information processing systems*, pp 6212–6221. <https://doi.org/10.5555/3454287.3454845>
 19. Štolba M, Urbanovská M, Fišer D, Komenda A (2019) Cost partitioning for multi-agent planning. In: *Proceedings of the 11th international conference on agents and artificial intelligence, ICAART*. <https://doi.org/10.5220/0007256600400049>
 20. Khambhaita H, Alami R (2017) Assessing the social criteria for human–robot collaborative navigation: a comparison of human-aware navigation planners. In: *2017 26th IEEE international symposium on robot and human interactive communication (RO-MAN)*, pp 1140–1145 (2017). <https://doi.org/10.1109/ROMAN.2017.8172447>. IEEE
 21. Khambhaita H, Alami R (2020) Viewing robot navigation in human environment as a cooperative activity. In: *Robotics research*, pp 285–300. Springer. https://doi.org/10.1007/978-3-030-28619-4_25
 22. Cai J, Du A, Liang X, Li S (2023) Prediction-based path planning for safe and efficient human–robot collaboration in construction via deep reinforcement learning. *J Comput Civ Eng* 37(1):04022046
 23. Chiou M, Hawes N, Stolkin R (2021) Mixed-initiative variable autonomy for remotely operated mobile robots. *ACM Trans Hum-Robot Interact (THRI)* 10(4):1–34. <https://doi.org/10.1145/3472206>
 24. Queralta JP, Taipalmaa J, Pullinen BC, Sarker VK, Gia TN, Tenhunen H, Gabbouj M, Raitoharju J, Westerlund T (2020) Collaborative multi-robot search and rescue: planning, coordination, perception, and active vision. *IEEE Access* 8:191617–191643. <https://doi.org/10.1109/ACCESS.2020.3030190>
 25. Islam MJ, Hong J, Sattar J (2019) Person-following by autonomous robots: a categorical overview. *Int J Robot Res* 38(14):1581–1618. <https://doi.org/10.1177/0278364919881683>
 26. Laplaza J, Garrell A, Moreno-Noguer F, Sanfeliu A (2022) Context and intention for 3d human motion prediction: experimentation and user study in handover tasks. In: *2022 31st IEEE international conference on robot and human interactive communication (RO-MAN)*, pp 630–635 (2022). <https://doi.org/10.1109/RO-MAN53752.2022.9900743>. IEEE
 27. Ferrer G, Zulueta AG, Cotarelo FH, Sanfeliu A (2017) Robot social-aware navigation framework to accompany people walking side-by-side. *Auton Robot* 41(4):775–793. <https://doi.org/10.1007/s10514-016-9584-y>
 28. Garrell A, Garza-Elizondo L, Villamizar M, Herrero F, Sanfeliu A (2017) Aerial social force model: a new framework to accompany people using autonomous flying robots. In: *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp 7011–7017. <https://doi.org/10.1109/IROS.2017.8206627>. IEEE
 29. Nguyen VT, Tran TD, Kuo I-H (2022) A model for determining natural pathways for side-by-side companion robots in passing pedestrian flows using dynamic density. In: *2022 31st IEEE international conference on robot and human interactive communication (RO-MAN)*, pp 1107–1114. <https://doi.org/10.1109/RO-MAN53752.2022.9900696>. IEEE
 30. Kohler D, Conley K (2011) Rosjava—an implementation of ROS in pure java with android support
 31. Quigley M, Conley K, Gerkey B, Faust J, Foote T, Leibs J, Wheeler R, Ng AY, et al (2009) Ros: an open-source robot operating system. In: *ICRA Workshop on open source software*, vol. 3, p 5. Kobe, Japan
 32. Fogg B (2009) A behavior model for persuasive design. In: *4th international conference on persuasive technology*, p 40. <https://doi.org/10.1145/1541948.1541999>
 33. Kitazawa K, Fujiyama T (2010) Pedestrian vision and collision avoidance behavior: investigation of the information process space of pedestrians using an eye tracker. In: *Pedestrian and evacuation dynamics 2008*, pp 95–108. Springer. https://doi.org/10.1007/978-3-642-04504-2_7
 34. Rios-Martinez J, Spalanzani A, Laugier C (2015) From proxemics theory to socially-aware navigation: a survey. *Int J Soc Robot* 7(2):137–153. <https://doi.org/10.1007/s12369-014-0251-1>
 35. Best G, Cliff OM, Patten T, Mettu RR, Fitch R (2019) Dec-mcts: decentralized planning for multi-robot active perception. *Int J Robot Res* 38(2–3):316–337. <https://doi.org/10.1177/0278364918755924>
 36. Dalmasso Blanch M (2020) Combining motion planning with social reward sources for collaborative human–robot navigation task design. Master’s thesis, Universitat Politècnica de Catalunya. <http://hdl.handle.net/2117/189415>
 37. Dalmasso M, Garrell A, Jiménez P, Sanfeliu A (2019) Human–robot collaborative navigation search using social reward sources. In: *Iberian robotics conference*, pp 84–95. https://doi.org/10.1007/978-3-030-36150-1_8
 38. Dalmasso M, Garrell A, Domínguez JE, Jiménez P, Sanfeliu A (2021) Human–robot collaborative multi-agent path planning using monte carlo tree search and social reward sources. In: *2021 IEEE international conference on robotics and automation (ICRA)*, pp 10133–10138. <https://doi.org/10.1109/ICRA48506.2021.9560995>. IEEE
 39. Domínguez-Vidal JE, Torres-Rodríguez IJ, Garrell A, Sanfeliu A (2021) User-friendly smartphone interface to share knowledge in human–robot collaborative search tasks. In: *2021 30th IEEE international conference on robot & human interactive communication (RO-MAN)*, pp 913–918. <https://doi.org/10.1109/RO-MAN50785.2021.9515379>. IEEE
 40. Sailer M, Hense JU, Mayr SK, Mandl H (2017) How Gamification motivates: an experimental study of the effects of specific game design elements on psychological need satisfaction. In: *Computers in human behavior*, pp 371–380. <https://doi.org/10.1016/j.chb.2016.12.033>
 41. Che Y, Okamura AM, Sadigh D (2020) Efficient and trustworthy social navigation via explicit and implicit robot-human communication. *IEEE Trans Robot* 36(3):692–707
 42. Bartneck C, Kulić D, Croft E, Zoghbi S (2009) Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *Int J Soc Robot* 1(1):71–81. <https://doi.org/10.1007/s12369-008-0001-3>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Marc Dalmasso is currently a PhD student in automatic control, computer vision and robotics at the Institut de Robòtica i Informàtica Industrial (CSIC-UPC), associated to the Universitat Politècnica de Catalunya (Barcelona-Tech). He previously received his B.Sc. in Computing Engineering and Industrial Engineering in 2016, his Master's Degree in Secondary Education Teaching in 2017 and his M.Sc. in Automatic Control and Robotics in 2020. His main contributions have been in human-robot collaboration in collaborative tasks while his interests focus on education, human-robot collaboration and negotiation and Roboethics.

J. E. Domínguez-Vidal is currently a PhD student in automatic control, computer vision and robotics at the Institut de Robòtica i Informàtica Industrial (CSIC-UPC), associated to the Universitat Politècnica de Catalunya - BarcelonaTech (UPC). He previously received his B.Sc. and M.Sc. in telecommunications engineering in 2016 and 2018 respectively, both from the University of Vigo, Galicia, Spain. His current interests are in the field of human-robot interaction, focusing on trying to understand the human intention, its implications and its applications to robotics based on concepts from psychology.

Iván J. Torres-Rodríguez is an electronics engineer from the University of La Laguna (ULL) with a double master's degree in industrial engineering and automatics and robotics from the Universitat Politècnica de Catalunya - BarcelonaTech (UPC). He has performed internships in the Institute of Astrophysics of the Canary Islands (IAC) and in SEAT. He has gained a collaboration grant from the Spanish ministry related to robotic arms. Apart from that he has been at the Institute of Robotics and Industrial Informatics (IRII) with a research grant for 2 years. And afterwards he has been as research assistant engineer also at IRII, where he has deepened into agricultural robotics with the CANOPIES project, computer vision techniques and machine learning.

Pablo Jiménez is Associate Researcher of the Spanish Research Council (CSIC) at the Institut de Robòtica i Informàtica Industrial (CSIC-UPC), where he currently holds the Department Head position. He studied Industrial Engineering at UPC and obtained his PhD at the Advanced Automation and Robotics program at Universitat Politècnica de Catalunya - BarcelonaTech (UPC). His main contributions have been in motion and task planning, and in robot learning. Recently his interests have focused on the robotic manipulation of deformable everyday objects like cloth, from the point of view of learning and planning, and with the objective of contributing to assistive Robotics in mind. Last but not least, he has become increasingly interested in issues related to Human-Robot Interaction in the context of Roboethics, the impact of Robotics on Human society from an ethical viewpoint.

Anais Garrell is an Assistant Professor at the Universitat Politècnica de Catalunya - BarcelonaTech (UPC) with a specialization in artificial intelligence applied to robotics; a researcher at the Institut de Robòtica i Informàtica Industrial, CSIC-UPC (IRI); and an expert in Human-Robot interaction. She received the BSEE degree of Mathematics from the Universitat de Barcelona, Spain, in 2006. She also obtained her PhD degree in Automatic, Control and Vision from the UPC, Barcelona, in 2013, and obtained the second prize of the best PhD Award in robotics in Spain. She has participated on diverse national and European projects and has worked with other research groups including the Robotics Institute of Carnegie Mellon University (US). Her research interests are focused on cooperative robotics, social robotics and human robot interaction. Dr. Garrell was the finalist of the "CEA Award for Female Automatic Talent" of Spain in 2022.

Alberto Sanfeliu received the Engineer and PhD degrees from the Universitat Politècnica de Catalunya - BarcelonaTech (UPC), Spain, in 1978 and 1982 respectively. He joined the faculty of UPC in 1981 and is full professor of Computational Sciences and Artificial Intelligence. He is the head of the research line Mobile Robotics of IRI and former coordinator of the research Artificial Vision and Intelligent System Group (VIS). He is former scientific director of the "Unity of Excellence Maria Maeztu" of IRI - Institut de Robòtica i Informàtica Industrial (CSIC-UPC). Prof. Sanfeliu is former director of IRI, former director of the UPC's Automatic Control Department, past president of AERFAI, (Spanish Association for Pattern Recognition) and past President of the Spanish Society for Research and Development in Robotics (SEIDROB). His research interest includes Intelligent Robotics, Human-Robot Interaction, Computer Vision and Pattern Recognition. He has several patents on quality control based on computer vision. He has authored books in pattern recognition and SLAM, and published more than 380 papers in international journals and conferences. He has led and participated in 64 R&D projects, 23 of them funded by the European Commission, one CONSOLIDER-INGENIO 2010, and he was the coordinator of the European project URUS (Ubiquitous Networking Robotics in Urban Areas). He has been member of editorial boards of top scientific journals in computer vision and pattern recognition and has been the General Co-Chairmen of the top international congresses of Computer Vision (ICCV'2011) and Pattern Recognition (ICPR'2000). He received the prize to the Technology given by the Generalitat de Catalonia and is Fellow of the International Association for Pattern Recognition.