

How to be Helpful?

Implementing Supportive Behaviors for Human-Robot Collaboration

Olivier Mangin, *Member, IEEE*, Alessandro Roncone, *Member, IEEE*, and Brian Scassellati, *Member, IEEE*

Abstract—The field of Human-Robot Collaboration (HRC) has seen a considerable amount of progress in the recent years. Although genuinely collaborative platforms are far from being deployed in real-world scenarios, advances in control and perception algorithms have progressively popularized robots in manufacturing settings, where they work side by side with human peers to achieve shared tasks. Unfortunately, little progress has been made toward the development of systems that are proactive in their collaboration, and autonomously take care of some of the chores that compose most of the collaboration tasks. In this work, we present a collaborative system capable of assisting the human partner with a variety of supportive behaviors in spite of its limited perceptual and manipulation capabilities and incomplete model of the task. Our framework leverages information from a high-level, hierarchical model of the task. The model, that is shared between the human and robot, enables transparent synchronization between the peers and understanding of each other's plan. More precisely, we derive a partially observable Markov model from the high-level task representation. We then use an online solver to compute a robot policy, that is robust to unexpected observations such as inaccuracies of perception, failures in object manipulations, as well as discovers hidden user preferences. We demonstrate that the system is capable of robustly providing support to the human in a furniture construction task.

I. INTRODUCTION

RECENT trends in advanced manufacturing and collaborative robotics are moving away from the traditional approach of building difficult-to-repurpose machines that work in isolation from humans. Rather, they are shifting focus to the development of mixed human-robot environments where robots are flexibly adaptable to the rapid changes of the modern manufacturing process and can safely and effectively inter-operate with humans. Unfortunately, albeit considerable progress in robot perception, manipulation and control has improved the robustness and dependability of such platforms, robots are still used as mere recipients of human instructions. That is, we still lack the breadth and depth to fully exploit the capabilities of the robot in general and the collaboration in particular. Human-robot collaborative systems are fundamentally unbalanced, with the bulk of the perceptual, cognitive and manipulation capabilities still pertaining to the human side of the system. For this reason, research in the field is shifting toward the deployment of systems that allow the human and




Fig. 1. The experimental setup, in which a human participant engages in a joint construction task with the Baxter Robot. In the picture, the robot is supporting its human partner by holding a leg of the table while the human is screwing. See Section IV-B for information about the task.

the robot to focus on the tasks for which they are best suited, and mutually assist each other when needed [30, 5, 15, 13].

In order for a robot to be proficient in the collaboration, we propose to focus our efforts on the implementation of platforms able to provide a variety of *supportive behaviors*. Our goal is for the robot to be capable of assisting the human *when the human needs support the most*, provided some degree of knowledge of the task at hand and sufficient information about the state of the system. Supportive behaviors such as handing over task components, providing tools, cleaning-up unused elements, holding a part during assembly happen to be extremely beneficial for the completion of the task, and are within the realm of possibilities for modern robotic platforms. Such behaviors may also cover information retrieval tasks such as lighting an area of a working space, providing execution time, or reminding parts of the task-plan. In such a scenario, it is not necessary for the robot to have exhaustive knowledge about the impending task, nor to be able to completely perceive the state of the world. Rather, for a system to exhibit effective support, partial observability about the state of the world and the internal state of the human partner (composed of her beliefs and intents) is sufficient. Importantly, we do not set forth to improve the physical capabilities of our robotic platform, but rather we focus on how to optimally exploit the limited control, perceptive and reasoning skills the robot has in order to best

O. Mangin, A. Roncone and B. Scassellati are with the Social Robotics Lab, Computer Science Department, Yale University, 51 Prospect St., New Haven, CT 06511, USA e-mail: name.surname@yale.edu

Manuscript received Month Day, Year; revised Month Day, Year.

support the human partner.

A key point for collaboration is that peers share a common understanding of the task, up to a certain degree [29, 30]. Because of the gap between the current cognitive capabilities of humans and robots it is both troublesome for the robot to reach the level of task understanding that the humans have, and impractical for the humans—let alone non-technical persons—to encode that task model in a way the robot can exploit. Therefore, we focus on basic models of the task, that are easy to describe for the human, yet contain enough information for the robot to be efficient and effective in its support. Notably, these shared task models can provide a substrate for human robot communication and thus foster transparent interactions between peers during task execution [17]. Importantly, the robot is only made aware of the part of the task that matters to it, which simplifies both task planning and plan execution. We do not aim to target robot controllers with exhaustive and exact knowledge—an approach that often falls short when outside of the context it has been designed for. Rather, we trade complete knowledge for adaptability, and optimal planning for good-enough support *by design*.

In this paper, we present a novel framework able to effectively empower a robot with supportive behaviors. We show, to our knowledge, the first practical implementation of an HRC application where the robot autonomously chooses to support the human when it deems it appropriate, and selects the right supportive action among the many it is provided with. Our approach aims to maximize the throughput of the mixed human-robot system by leveraging the superior perceptual and manipulative skills of the human partner, while entrusting the robot with the role of autonomous helper, for which it is best suited for. Additionally, we intend to systematically employ the human partner to retrieve the knowledge the robot is not able to gather by itself, in order to improve its estimation of the state of the world and the task completion. Finally, our system is also able to adapt to user’s preferences in terms of when and where to provide supportive behaviors. In our experiments, we further demonstrate how the system is able to dynamically comply with users’ preferences *during execution*. We validate the proposed framework in joint construction tasks that aim to simulate the manufacturing process typical of Small and Medium Enterprises (SMEs), where features such as reconfigurability and ease of deployment are paramount.

In the following sections, we introduce the reader to the state of the art and related works in the field (Section II). Then, we detail the proposed approach, focusing on how it differentiates from relevant research in the topic (Section III). The experimental setup and the experiment design are presented in Section IV, followed by the Results (Section V) and Conclusions (Section VI).

II. BACKGROUND AND RELATED WORK

In this work, we capitalize on past research in the field of high-level task reasoning and representation. As detailed in Section III, the core contribution of this paper is a system able to convert hierarchical task models into low-level planners capable of being executed by the robot. Contrarily to more

traditional techniques that leverage full observability in the context of HRC applications (e.g. [21, 34]), our system deliberately optimizes its actions based on the interaction dynamics between the human and the robot. We explicitly account for uncertainty in the state of the world (e.g. task progression, availability of objects in the workspace) as well as in the state of the human partner (i.e. her beliefs, intents, and preferences). To this end, we employ a partially observable Markov decision process (POMDP) framework that plans optimal actions in the belief space.

To some extent, this approach builds on top of results in the field of task and motion planning (TAMP, see e.g. [23, 22, 24]). Indeed, similarly to Kaelbling and Lozano-Pérez [22] we find approximate solutions to large POMDP problems through planning in belief space combined with just-in-time re-planning. Our work differs from traditional TAMP approaches in a number of ways: i) the hierarchical nature of the task is not explicitly dealt with in the POMDP model, but rather at a higher level of abstraction (that of the task representation, cf. Section III-A), which reduces complexity at planning stage; ii) we encapsulate the complexity relative to physically interacting with the environment away from the POMDP model, which results in broader applicability and ease of deployment if compared with standard TAMP methods; iii) most notably, our domain of application presents fundamental differences with that targeted by TAMP techniques. That is, we propose to *handle uncertainty in the human-robot interaction*, rather than in the physical interaction between the robot and the environment. This latter point is worth elaborating on, since under this paradigm there is no shared consensus on how to model uncertainty about human’s beliefs and intents in general, and the collaboration in particular.

Planning techniques can enable human robot collaboration when a precise model of the task is known, and might adapt to hidden user preferences as demonstrated by [35]. Similarly, partially observable models can provide robustness to unpredicted events and account for unobservable states. Of particular note is the work by Gopalan and Tellex [9] which, similarly to the approach presented in this paper, uses a POMDP to model a collaborative task. Indeed, POMDPs and similar models (e.g. MOMDPs) have been shown to improve robot assistance [16] and team efficiency [26] in related works. Such models of the task are however generally expensive to build and require advanced technical knowledge. Hence, a significant body of work in the fields of human-robot collaboration and physical human-robot interaction focuses on how to best take over the human partner by learning parts of the task that are burdensome in terms of physical safety or cognitive load. Under this perspective, the majority of the research in the field has focused on frameworks for learning new skills from human demonstration (LfD, [2]), efficiently learn or model task representations [14, 34, 8, 19], or interpreting the human partner’s actions and social signals [10].

No matter how efficient such models are at exhibiting the intended behavior, they are often limited to simple tasks and are not transparent to the human peer. Indeed, evidences from the study of human-human interactions have demonstrated the importance of sharing mental task models to improve

the efficiency of the collaboration [30]. Similarly, studies on human-robot interactions show that an autonomous robot with a model of the task shared with a human peer can decrease the idle time for the human during the collaboration [31]. Without enabling the robot to learn the task, other approaches have demonstrated the essential capability for collaborative robots to dynamically adapt their plans with respect to the task in order to accommodate for human’s actions or unforeseen events [17]. Likewise, rich tasks models can also enable the optimization of the decision with respect to extrinsic metrics such as risk on the human [18] or completion time [28].

Our paper is positioned within this growing body of work related to task representations in HRC. Unfortunately, little attention has been given to the issue of explicitly tackling the problem of effectively supporting the human partner. To our knowledge, Hayes and Scassellati [13] is the only work that goes in this direction. It presents an algorithm to generate supportive behaviors during collaborative activity, although its results in simulation fall short in terms of providing practical demonstrations of the technique. On the other side of the spectrum, a number of works cited above achieve to a certain amount supportive behaviors without explicitly targeting them [18, 31, 9, 34]. A limitation of these approaches is that, as mentioned previously, they rely on exact task knowledge that is not always available for complex tasks in practical applications.

III. METHOD

This work capitalizes on previous research by the authors. In Roncone et al. [28], we demonstrated an automated technique able to derive robot-executable policies from human-readable task models. We then exploited this framework in the context of role assignment: our system was effective in efficiently negotiating allocation of a specific subtask to either the human or the robot during a collaborative assembly. In this work, we expand this approach toward the more general problem of optimally providing support to the human. Similarly to our previous work [28], we employ hierarchical representations of the task at a level of abstraction suitable to naive human participants and understandable by the robot. The model of the task is provided a priori, although other studies have shown that it is possible to learn task models from human demonstrations [6, 12]. We then convert this task representation to a robot policy by leveraging the flexibility of POMDP models. This allows the robot to plan under uncertainty and explicitly reason at a high level of abstraction.

It is worth noting that exploiting adaptive planning from POMDPs has been already demonstrated in the context of human-robot collaboration (e.g. [9, 26]). Planning under uncertainty is indeed a major requirement for robots to interact with people. Near-future robotic platforms are most likely to operate in highly unstructured environments, for which even state of the art perception systems are not going to provide full observability or exact estimations. In this work, we push this idea to its limits, inasmuch as we constrain our framework to the condition of being nearly *blind*. That is, the robot is not able to directly observe neither the state of the world (task

progression, object locations, etc.), nor the state of the partner (intents, preferences, etc). This allows to investigate mechanisms of coordination through communication and physical interaction in the environment. We achieve that by expanding the technique introduced in [28], in that we leverage a high-level task model to *automatically generate* the lower-level POMDP. We then demonstrate how the resulting policy is successful in providing support during realistic human robot collaborations.

A. Hierarchical Task Models (HTM)

Hierarchical structures form an appealing framework for high-level task representations; of particular interest is their capability to enable reuse of components over different tasks. Additionally, their level of abstraction is usually close to human intuition: this facilitates human-robot communication about task execution [28].

Figure 2 depicts some example representations for real-world construction tasks. Similarly to Hayes and Scassellati [14], we consider HTMs built from primitive actions with operators that combine them into what we refer to as *sub-tasks* of increasing abstraction. In this work, we assume that information about a set of primitive actions is already available to the robot, and we represent complex tasks on top of this action vocabulary. We assume also that the robot has basic knowledge pertaining these primitive actions. This can range from knowing the type of tools and parts needed to perform an action, to being aware of the fact that supporting the human through holding a part may be beneficial during complex executions. In our previous work [28], we extended the CC-HTM representation introduced in [14] with the introduction of a new *alternative* operator (\vee in Fig. 2). It adjoins the *sequential* (\rightarrow) and *parallel* (\parallel) operators. This set of operators proves suitable to capture the complexity of the collaboration, as well as the constraints of a task execution. For example, the parallel operator allows for the two peers to perform two disjoint subtasks at the same time; conversely, the sequential operator constrains them to a specific sequence of execution.

Thanks to their simplicity, HTM models can conveniently be drafted by non-expert workers and remain intuitive to understand. Differently from traditional TAMP approaches [21, 22], their high-level of abstraction also enables decoupling of the task planning component from the robot control element. This increases flexibility, in that the robot just needs to be equipped with some motor primitives to match the atomic actions that compose the HTM. From then on, the same library of motor primitives can be used to repurpose the robot to a new task. Furthermore, it becomes easy to port the same task to a new platform with comparable motor and perception skills.

One of the main limitations of HTM-like approaches to HRC is that it is unlikely for the designer to be able to encode the totality of the information about the actual components of the task. Although this poses limits to the breadth of applicability of such techniques, we argue that, for a robot to provide effective support, perfect knowledge about task execution is not needed in the first place. A partial HTM, and specifically one that conveniently encodes only the information

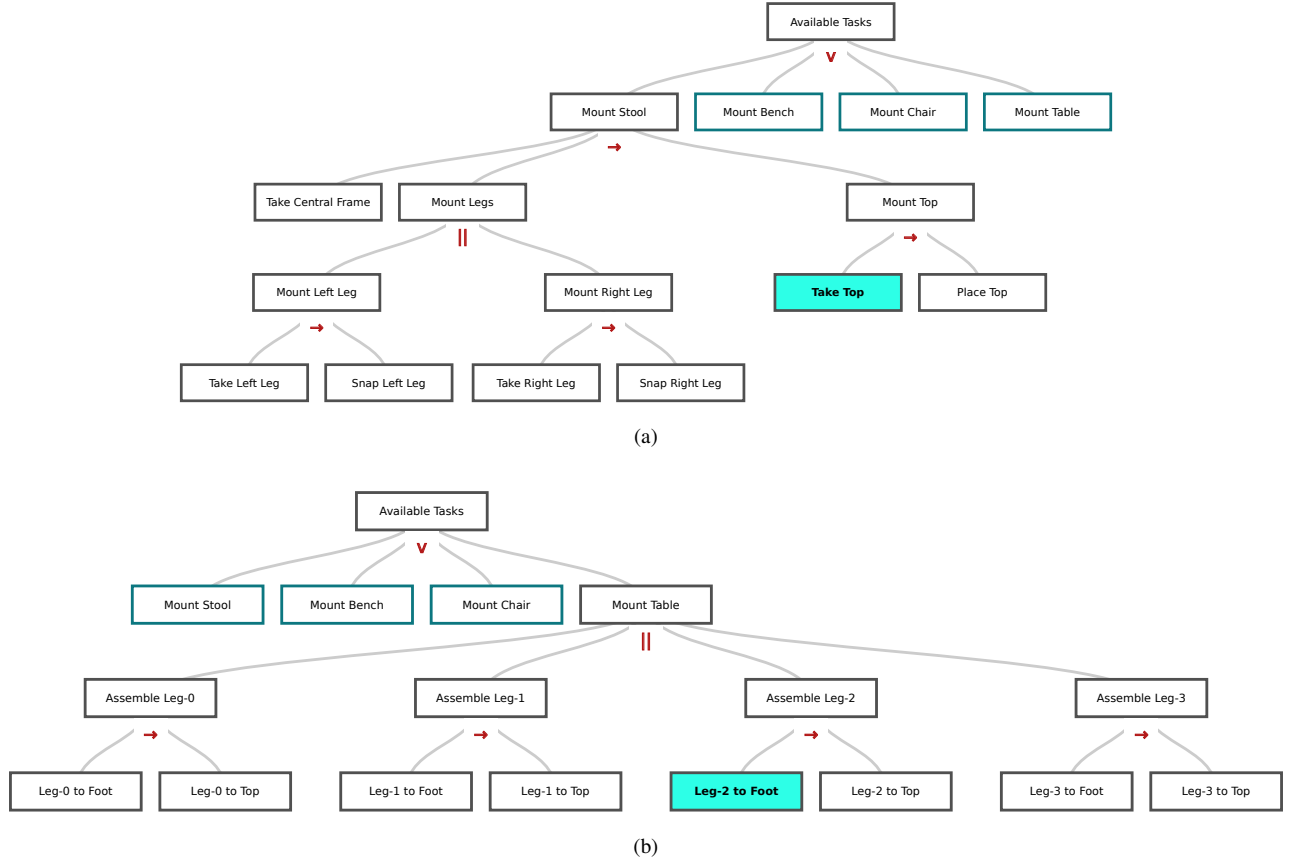


Fig. 2. Example HTMs for human-robot assembly tasks. The user can retrieve information on the task execution, and inquire the system about task progression that will be highlighted by the robot during execution (cyan block in picture). The operator type between subtasks (\rightarrow for sequential tasks, $||$ for parallel tasks, \vee for alternative tasks) is also available for the user to inspect. **a)** HTM used in Roncone et al. [28]. **b)** HTM used in Section V-B. See text for details.

that matters to the robot, is sufficient for it to operate and interact with the human. For example, the robot does not need to know how to perform a screwing action, nor how to perceive progression of the human screwing. What matters for a supportive robot is what objects are needed to complete said action, and that the human may be facilitated if the robot holds the part steadily. As discussed in Section III-B, our POMDP model complements partial knowledge about the task and the state of the world through interaction with the human partner. It can for example supplant lack of perception about subtask progression by asking the human when the subtask is completed, or by directly moving on to the next subtask if its likelihood of subtask completion is high enough.

B. Partially Observable Markov Decision Processes (POMDP)

In this work, we use POMDPs to formulate the decision problem that the robot faces, given a task to solve collaboratively with a human and represented as an HTM as explained in Section III-A. POMDPs are a generalization of Markov decision processes (MDPs), where there is only partial observability of the state of the process. This important relaxation of what defines an MDP allows for a significant gain in flexibility. It is particularly relevant to model-imperfect perception and hidden states such as user preferences. We use such an approach to optimize the robot actions despite

incomplete knowledge of the task and uncertainty regarding the dynamics of the collaboration.

More precisely, a POMDP is defined by a 7-tuple $(S, A, \Omega, T, O, R, \gamma)$, where S is a set of states, A is a set of actions, T is a set of state transition probabilities, $R : S \times A \rightarrow \mathbb{R}$ is the reward or cost function, Ω is a set of observations, O is a distribution of observation probabilities, and $\gamma \in [0, 1]$ is the discount factor. Similarly to a MDP, at any given time the system lies in a specific state $s \in S$, which in the case of POMDPs is not directly observable. The agent's action $a \in A$ triggers a state transition to state $s' \in S$ with probability $T(s' | s, a)$ and an observation $o \in \Omega$ with probability $O(o | s', a)$ that depends on the new state s' . Finally, the agent gets a reward $r \in R$ for taking the action a while in state s . In case of POMDPs, the agent's policy is defined on a probability distribution over states b , called the belief state, which accounts for the fact that the agent has no direct access to the real state s . The goal is for the POMDP solver to find a policy $\pi(b) : b \rightarrow a$ that maximizes the future discounted rewards over a possibly infinite horizon: $E[\sum_{t=0}^{\infty} \gamma^t r_t]$. Interestingly, actions that do not change the underlying state of the system but only the belief state are also valuable in this context. This proves particularly beneficial for human collaboration, since information-gathering actions belong to this paradigm. For example, this intuition can be used to model communicative actions that trigger observations to

disambiguate uncertainty, or to favor low-entropy beliefs with small uncertainty for both the human and the robot. In reality, the belief state is usually very large and continuous; we use a policy that is defined on the history of previous actions and observations that we denote by $h \in H$. Please refer to Section III-D for more information on how we compute a robot policy from POMDP models.

C. Restricted Model (RM)

We propose an automated technique able to transform task-level HTMs into low-level robot policies through POMDPs. To this end, we convert each primitive subtask (that is, each leaf composing the HTM in Figs. 2a and 2b) into a small, modular POMDP, which we call a restricted model (RM [32]). Hence, each RM is mostly independent from the rest of the problem and can be studied in isolation. Differently from [32], the RMs are composed at a later stage, according to the HTM structure, and the problem is solved in its entirety. This approach benefits from the modularity of the HTM representation, without the typical sub-optimality of policies that do not consider the full problem.

Figure 3 depicts the RM developed in this work. As mentioned in Section III-A, its action space corresponds to the set of motor primitives available to the robot. For the purposes of this work, we consider the following supportive actions: i) *wait* for the human to complete a subtask; ii) *hold* an object to provide support to the human; iii) *bring object* (e.g. constituent parts, small parts, buckets, or tools) onto the workspace; iv) *clean-up object* from the workspace when not needed anymore. These motor primitives are implemented as independent controllers with their own logic: for example, the *wait* controller exploits communication in order to ask the human when the current operation has been completed before moving on to the new subtask. Modularity is employed in order to derive a distinct action for each object involved.

The set of possible observations is limited to a *none* observation (the default), plus a set of *error* observations returned by either the robot itself (e.g. *object-not-found*, *kinematic-error*) or the human partner (e.g. *wrong-action*). Forcing the system to deal with a limited set of observation is intentional, and resorts to the intuition that, when the robot's execution is correct, the human partner should not be concerned about reinforcing this with positive feedback—which is time consuming and cognitively taxing. Rather, feedback from the human partner should come either if explicitly requested by the robot—to disambiguate uncertainty, or if the robot's decision making is wrong—and negative feedback can be used to correct its course of actions.

Lastly, the state space S is composed of a set of factored states. It is conceptually divided into three sub-spaces: i) an *HTM-related state space* Q pertains to task progression, and is derived directly from the HTM representation. Each of the subtasks in the HTM (i.e. each of the leaves) is assigned an unique state $q \in Q$; additionally, one final state \hat{q} is associated with a virtual operation that requires the robot to cleanup the workspace. ii) a *subtask-related state space* is defined alongside the controllers that perform each subtask;

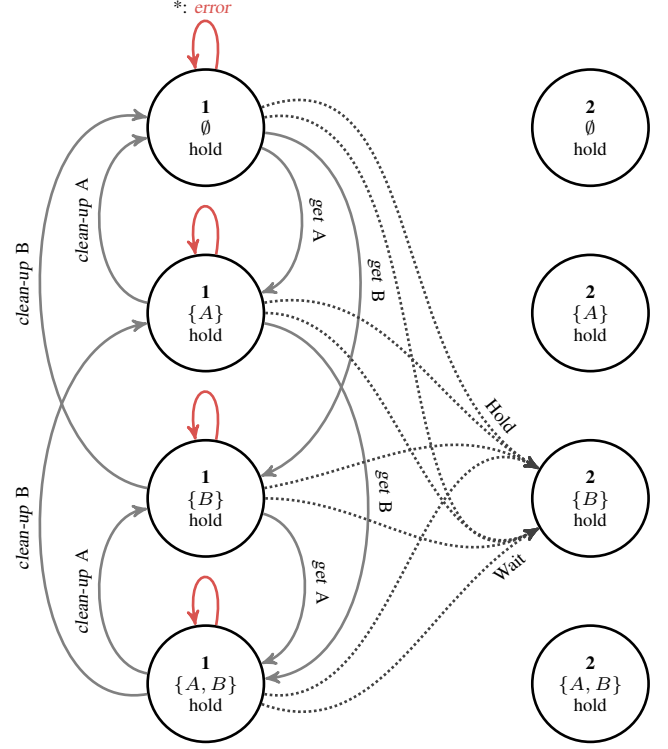


Fig. 3. Simplified representation of the restricted model (RM) used in this work. The figure represents an RM associated with a subtask 1 whose successor in the HTM is subtask 2. For the sake of simplicity, the figure only represents actions that are taken starting from subtask 1. We assume that only two objects 'A' and 'B' are available and that 'A' is consumed by the subtask (like a part of the assembly would be) while 'B' is a tool used during the task (e.g. a screwdriver). Each node represents a state, that is a factorization of the HTM subtask, each possible combination of objects on the workspace, and the user preference regarding the *hold* supportive action. Full connections in the graph represent successful transitions for the actions *get* and *clean-up* applied to objects 'A' and 'B' (that lead to a *none* observation). When taken from other states (e.g. bringing 'A' which is already on the workspace), the action would fail, with an *error* observation, and the state would not change. These cases are represented by the red connections. Finally, the dotted connections represent the *hold* and *wait* actions that, from any of the represented states, lead to a transition to the state of the next subtask for which object 'B' only is present (i.e. the tool). This means that the transition occurs even if the robot failed to bring all the required tools and parts: we assume here that the human would be able compensate for robot failures. The reward would however be maximal in the transition from state with $\{A, B\}$. To simplify the figure, we omitted the states corresponding to the *no-hold* preference. The graph for the *no-hold* preference is nearly identical except for the fact that the *hold* action fails from these states and hence *wait* is the only action to transition to the next subtask.

it is composed of information relating parts and tools (e.g. if tools are present in the workspace, or if parts have been 'consumed'). iii) a *human-related state space* encompasses information related to human preferences, beliefs and intents. As shown in Section IV-C, in this work we demonstrate how the proposed approach is able to adapt to user preferences regarding the *hold* supportive action.

It is worth noting how the size of the state space S grows exponentially with the number of preferences and objects, as detailed in Section III-D. In order to account for scalability of the method, we define a *generative POMDP model* that circumvents the issue of explicitly defining the full transition matrix T . Instead, as detailed in Fig. 3, we generate it as

TABLE I

REWARDS USED TO TRAIN THE POLICY ON THE POMDP MODEL DERIVED FROM THE HTM. INSTEAD OF A TABLE OF REWARDS FOR ALL STATE TRANSITIONS, ACTIONS, AND OBSERVATIONS, WE PRESENT THE REWARDS AS TRIGGERED BY EVENTS THAT CAN BE CUMULATED. FOR EXAMPLE, IF A FINAL STATE IS REACHED THROUGH A WAIT ACTION, BUT THE SCREWDRIVER IS STILL ON THE WORKSPACE, A REWARD OF 85 IS OBTAINED.

Event	Reward
Final state reached	100
Subtask transition	10
Missing tool or part on state transition	15
Uncleaned object on final state	15
User preference is honored	10
Hold action taken	-2
Wait action taken	0
Other action taken	-1

follows. Each action affecting an object changes the state representing its presence in the workspace: for example, bringing a screwdriver makes it available on the workspace with high probability. The *wait* action and eventually the *hold* action trigger the transition from one HTM leaf to the next (according to their order in the sequence, the final operations leading to a transition to the special state \hat{q}). This mechanism enforces that transitions between HTM states are transparently synchronized with the human. *Hold* only triggers the transition from states that have the preference for holding and fails otherwise. Additionally, the transition from one leaf to the next erases from the state representation all the objects that have been ‘consumed’ by the subtask (typically, the parts that have been used). The initial state is sampled by starting at the initial subtask $q_0 \in \mathcal{Q}$; the workspace is assumed to be free of objects, and the human preferences are randomly set.

Interestingly, the design choice of limiting the perception of the state of the world naturally conforms to the statistical nature of a POMDP approach. That is, adding uncertainty in the model makes it ultimately more robust to actual uncertainty in the collaborative interaction. Without loss of generality, the RM can allow for unexpected transitions in order to account for actions of the human that are not observed by the robot—e.g. when the human partner fetches a required component by herself, unexpected failures, or missing objects. To model this, and to avoid the robot to be stalled in a wrong belief, we introduce low probability random transitions between all the state features in S .

Finally, rewards are provided in the following cases: i) when the robot proposes to hold, and the human has preference for holding; ii) when there is a transition between a subtasks and its successor; ii) at completion of the full task. Additionally, each action taken by the robot has an intrinsic cost, and a negative reward is also given when the human has to bring or clean an object which was not taken care of by the robot. Table I provides a summary of the rewards used in the experiment from Section V-B.

D. POMDP Planner

In this work, we implement a planner based on POMCP [33], which is able to plan from generative models

and can handle very large state spaces. This is achieved through Monte-Carlo estimation of the beliefs by using a set of particles to represent each belief. A policy is learned based on all the visited histories, which constrains exploration to feasible states only. Using particles for belief representation and Monte-Carlo techniques for value estimation addresses the issue of the belief space being too large to be explicitly represented in its entirety. In a realistic HRC domain—such as those detailed in Section V—there are typically thousands of states and tens of actions, but the amount of plausible states at any time is limited. Hence, representing the $|\mathcal{S}|$ -dimensional belief is not feasible, but at the same time, despite large state spaces, beliefs are sparse. This is well represented through sets of particles, that naturally conform to sparse representations. Further, this approach only requires a generative model of the transitions instead of representing the full transition matrix, whose dimension is $|\mathcal{S}|^2 \times |\mathcal{A}| \times |\mathcal{O}|$.

The computational complexity of the planning is bound to an exploration-exploitation trade-off. Namely, the planner explores a tree whose branching factor is equal to the number of actions multiplied by the number of observations. Therefore, an important parameter to control the complexity of the problem is the horizon of the exploration. More precisely, we define the horizon either in terms of the number of transitions or as the number of *HTM subtasks* that the exploration accounts for. Such optimizations result in locally optimal decisions for what concerns a fixed number of subtasks.


In order to limit computational complexity, it is also possible to remove part of the randomness introduced in Section III-C for what concerns the transitions and the observations—although it still occurs in the belief transitions that hence still represent all the possible hypothesis. By doing so, we prevent the planner from exploring feasible but rare events. In case these rare events occur during live interaction with the human, the online component of the planner is able to recompute a new policy. This makes the algorithm robust to unexpected events without penalizing the exploration. That is, we artificially simplify the model of the interaction at the offline planning step, but we then compensate for its imperfections online—during task execution.

IV. IMPLEMENTATION


A. Experimental Setup

The experimental evaluation is carried out on a Baxter Research Robot (cf. Fig. 1), using the Robot Operating System (ROS [27]). As mentioned in Section III, even though we do not concern with improving the physical capabilities of the platform, we leverage the state of the art in robot perception and control in order to build up a set of basic capabilities for the robot to effectively support its human partner. The resulting framework, originally developed in Roncone et al. [28], exposes a library of high level actions, that are the only interface through which the POMDP planner can send information to—and retrieve information from—the Baxter system and the experimental setup.

The system presented in [28] provides multiple, redundant communication channels to interact with the human partner



(a) Left end effector's camera view, which is partly obscured by the vacuum gripper (top part of picture).



(b) Right end effector's camera view, which is partly obscured by the parallel gripper (top part of picture).

Fig. 4. Snapshot of the camera streams from the left (Fig. 4a) and right (Fig. 4b) end effectors. The left arm uses a fiducial marker tracking system based on [7], whereas the right arm implements the HSV-based 3D reconstruction software detailed in Section IV-A1. See Section IV-B below for a description of the objects composing the construction task.

‘on human terms’[4]. Among the list of available layers, for the purposes of this work we employ: i) a *Text-to-Speech (TTS)* channel, used to verbally interact with the human; ii) a *Feedback* channel, shown in the robot’s head display, which provides feedback about its internal states and intents (see Fig. 1); iii) an *Error* channel that allows the human to send error messages to the robot, and is triggered by pressing one of the buttons on the robot’s end effectors. In addition to this, a fourth channel has been implemented, in the form of a *Speech-to-Text (STT)* system able to convert human sentences into robot-readable commands. It employs the Google Cloud STT API [1] combined with a text parser that relies on a dictionary shared in advance with the human participant.

Both arms are able to perform precise, closed loop visual servoing tasks thanks to a pair of cameras their end effectors are equipped with. The left arm is equipped with a vacuum gripper, able to pick up flat surfaces with constant texture, whereas the right end-effector is a parallel electric gripper capable of performing more complex grasping tasks. We maximize the usage of both arms by leveraging their respective embodiments: to this end, two different perceptual systems have been employed (cf. Fig. 4). The perception system for

the left arm (Fig. 4a) is provided by ARuco [7], a library capable of generating and detecting fiducial markers that are particularly suitable for being positioned on flat surfaces. For what concerns the right arm, a custom color-based pose estimation algorithm has been implemented. It is detailed in the following section.


1) *6D object reconstruction from single view*: We consider here the scenario in which the end-effector is vertically placed on top of the pool of objects, and the objects are in the field of view of the camera. In order to be able to grasp a variety of objects with the parallel gripper installed on the Baxter’s right arm (see Fig. 4b), the following two steps are to be performed: objects need to be firstly detected in the camera view, and then their position and orientation has to be reconstructed in the 3D operational space of the robot. For what concerns the former, a number of different computer vision techniques can be employed. In this work, we utilize a Hue-Saturation-Value (HSV) color segmentation algorithm: that is, each object is detected thanks to its color in the HSV color space, and its bounding box is stored for later use. After an object has been detected, it is necessary to estimate its 3D pose in the world reference frame. We assume here that the object’s physical sizes (width and height) are known, and that the matrices of intrinsic and extrinsic parameters K and $[R|T]$ are available. Notably, whilst K can be estimated via a prior camera calibration step, the extrinsic parameters are computed thanks to the robot’s kinematics and the knowledge of the current joint configuration. In this context, a standard perspective transformation can be applied in order to estimate the 3D position of a point $P_w = [X \ Y \ Z \ 1]^T$ in the world reference frame from its corresponding image point $p_c = [u \ v \ 1]^T$ in the camera reference frame. The following equation holds:

$$s \ p_c = K [R|T] \ P_w, \quad (1)$$


where s is a scale factor¹. The perspective transformation equation is then applied to estimate the pose of the mass center of the object by iteratively minimizing the reprojection error of its corners via a Levenberg-Marquardt algorithm [25], using the OpenCV computer vision library [3].

The technique proposed here is subject to a number of estimation and computational errors, in particular if the distance between the camera and the desired object is significant. In spite of that, we capitalize on the fact that the algorithm is employed in a visual servoing setup, in which the robot refines its estimation the closer it gets to the object. That is, even if the initial pose reconstruction may be defective, it is continuously updated with a frequency of 30 Hz and refined until the end effector reaches the object. The authors acknowledge that more advanced 3D reconstruction techniques could be used, such as exploiting a depth sensing camera properly calibrated with respect to the robot—e.g. [20]. The main advantage of the proposed solution is however to employ a compact, self-contained estimation step that does not rely on external equipment or burdensome calibration. We consider

¹ It is out of the scope of this work to present the perspective transformation problem in detail. Please refer to Hartley and Zisserman [11] for more information on the topic.



(a) Parts composing the table.



(b) Completed table.

Fig. 5. **a)** The table building task is composed of one plywood tabletop, four dowels that act as legs, four brackets (top 3D printed objects in figure) and four feet (bottom 3D printed objects in figure). A total of 16 screws are needed to secure parts together. Both the tabletop and the legs have been pre-drilled to facilitate assembly. **b)** The table after the construction task is completed. The only tool required for the assembly is a screwdriver (red object in bottom of Fig. 5b).

this an important asset of our approach, that facilitates re-use and applicability to novel domains.

B. Experiment Design

As detailed in Sections I and IV-A, we perform our experiments in a collaborative scenario where human participants engage in a construction task with the Baxter Robot (see Fig. 1). The collaborative task the two peers are engaged with is the joint construction of a miniaturized table (cf. Fig. 5). It is composed of five structural elements—the tabletop and four legs—and eight custom 3D-printed linkages—four brackets secure the legs to the tabletop, whereas four feet are used to stabilize the structure. A total of 16 screws are required for the assembly. A screwdriver is the only tool needed to build the table, which has an approximate size of $30 \times 21 \times 15$ cm when completed.

Particular attention has been placed in the conceptualization of the task. The main goal was to tailor the design of the table to the typical constraints human-robot collaboration

experiments present. We purposely aim for: i) ease of reuse; ii) ease of retrieval of the constituent parts; iii) scalability; iv) proximity with real-world HRC applications, that are typically characterized by a combination of complex actions to be performed by the human partner (which often involve the use of tools), and simpler tasks the robot is usually assigned to. The choices taken at the design stage allow us to comply with these requirements: the brackets used in this experiment belong to a larger library of linkages that has been made available online², whereas the table shown in Fig. 5b is only one of the many designs allowed by our solution. We plan on leveraging these features in future works, and to promote this design as a model set to benchmark human-robot collaboration.

For the purposes of this work, the two partners have distinct, non overlapping roles: the human (hereinafter referred to as the *builder*) is in charge of performing actions that require fine manipulation skills (e.g. screwing) or complex perception capabilities (such as inserting the top of the table onto a bracket); the robot (also called the *helper*) is instructed to back the builder with the supportive actions described below. Importantly, the flexibility of the POMDP planner allows for a certain slack in terms of role assignment and task allocation by design. As detailed in Section III, the planner is automatically able to comply with unlikely states of the system, and to re-plan accordingly. As a practical consequence to this, we are in the position of allowing the human participant to take charge of some supportive actions if he so chooses. That is, we disclose to the builder that she is allowed to retrieve parts and tools for herself, even though we do not enforce that on the user—nor we convey that it is part of her duties as a participant of the experiment.

The robot helper is provided with a set of basic capabilities encapsulated into a library of high-level actions. The supportive actions it has been instructed to perform are the following: i) retrieve parts (e.g. tabletop, screws, or legs); ii) retrieve the tool (namely, the screwdriver); iii) cleanup the workplace from the objects that are not going to be needed in the future; iv) hold structural parts in order to facilitate the builder’s actions. To comply with the Baxter’s limited manipulation capabilities, we positioned the smaller components (i.e. the screws and the 3D printed objects) in apposite boxes, to be picked up by the parallel gripper—see *screws_box* and *brackets_box* in Fig. 4b. Similarly, the legs of the table have been equipped with a specific support in order to be picked up by the vacuum gripper on the left arm (cf. Fig. 4a).

It is worth noting how this specific task is particularly advantageous for the purposes of this work thanks to i) its simplicity and ii) the need for the human participant to perform the same actions multiple times. We deliberately designed a task that does not require any particular skill from the builder, while being easy to understand and remember. Although it may be tedious for the user, the need of performing multiple actions of the same type is beneficial in terms of showcasing the online user adaptation capabilities introduced in Section I. As detailed in Section V below, one of the assets of the

² scazlab.github.io/HRC-model-set hosts CAD models, specifications for 3D printing, tutorials for assembling example designs, and reference links for purchasing parts.

proposed system is to be able to abide by the builder’s preferences: in such a scenario, the robot is able to receive eventual negative feedback in case of wrong action, replan accordingly, and exhibit the effects of such replanning *within the same task execution*, i.e. without having to perform a new task from scratch.

C. Experimental Evaluation

We demonstrate the proposed system in a live interaction with human participants. The robot is in charge of backing the user with the right supportive action at the right moment by virtue of a partial observation of the state of the world, the complete knowledge of the task execution plan, and the HTM-to-POMDP planner presented in Section III. We devised two distinct experimental conditions, in a within-subjects design. For all the conditions, the skill set and capabilities of the robot do not vary, but the user preferences are explicitly altered, unbeknownst to the robot. That is, the robot is exposed to a change in the state of the system—composed of the world plus the human—that it could not observe, but needs to deduce either by actively gathering information from the builder, or by building upon feedback coming from her. In the following sections, the two experimental conditions are detailed. Please refer to Section V for a comparative evaluation.

1) *Condition A - full support*: In this scenario, the robot expects to support the human to the best of its capabilities, that is by performing all the actions it is allowed to. Firstly, the builder is introduced to the platform and the construction task. The experimenter then proceeds to illustrate the Baxter’s capabilities (i.e. providing parts, retrieving tools, holding objects and cleaning up the workspace) and the interaction channels the human is supposed to employ during task execution. Next, the experimenter communicates to the user that the robot is supposed to perform all the supportive actions by itself, but also that the participant is free to take charge of some actions if she so chooses or if the robot fails. No information is given in terms of what to expect from the robot, or how the human-robot interaction is supposed to occur.

2) *Condition B - adaptation to user preferences (no holding actions required)*: This condition involves the same interaction between the human and the robot as Condition A. As detailed in Section IV-C, the independent variable we tweak in our within-subjects experiments is the user preference for what concerns the support that the human participants expects from the robot. In this scenario, the human worker is told to prefer not to have parts held by the robot while screwing. Since the robot is unaware of this, it may still perform the holding action even if not required. In case this happens, the human is instructed to negatively reward the robot by sending an error signal to the Baxter.

V. RESULTS

The framework detailed in Sections III and IV-A has been released under the open-source LGPLv2.1 license, and is freely available on GitHub³. A number of C++ based ROS

packages has been made available for robot-related software, whereas the planner has been encapsulated into a ROS-independent Python package. In the following sections, we evaluate the proposed approach. Firstly, we perform a series of off-line experiments to assess if the proposed model can derive effective policies against a variety of tasks and experimental conditions. We show how our method outperforms ad-hoc policies on simulated interactions from these models (Section V-A). Lastly, we validate our system in a live interaction with human participants. We demonstrate how effective task policies can be computed that enable supportive behaviors during collaborative assembly tasks (Section V-B).

A. Off-line evaluation of the method

In this section, we present a quantitative evaluation of the proposed approach during off-line explorations. To this end, we focus on the two most important aspects that are involved in the design of effective human-robot interactions: i) the flexibility of our method against a variety of task structures (Section V-A1) and ii) its adaptability to custom user preferences (Section V-A2).

1) *Task structures*: In order to demonstrate how the proposed method is able to provide effective support to a human partner in real-world collaborative tasks, we evaluate it on three task models, derived from distinct HTMs. The HTMs differ in the number of subtasks to solve, and in the type of relational operators between subtasks. They hence illustrate how we can derive policies from various task models. All task models in this section are characterized by primitive subtasks that require a combination of: i) a set of tools that the robot needs to initially bring and then clean at the end of the task; ii) a shared supportive action ‘a’; iii) another supportive action, that can be either ‘b’ or ‘c’. Let denote ‘B’ the subtask that involves supportive actions ‘a’ and ‘b’ and ‘C’ the one involving ‘a’ and ‘c’. The first HTM, denoted as *sequential* task, consists of a sequence of 20 subtasks ‘B’. The second task model is denoted as *uniform*; it consists of an alternative between 24 subtasks, each composed of a sequence of four subtasks, each of type ‘B’ or ‘C’. In other words, for each episode the current task is randomly chosen among any sequence of three subtask of type ‘B’ or ‘C’. The last HTM, denoted as *alternative* task, is an alternative between only 4 sequences of four subtasks: ‘BCCC’, ‘BBBB’, ‘CBBC’, and ‘CCBC’. It thus introduces a dependency between the successive required actions.

We compare the performance of the proposed approach against two hand-coded policies. A *random* policy initially brings all the required tools (which is always a successful strategy); it then takes action ‘a’, and after that it randomly chooses between action ‘b’ and ‘c’ until one succeeds. It finishes by cleaning the workspace. When observing a failure (except on ‘b’ and ‘c’), the policy simply repeats the last action. The *repeat* policy is instead designed for the *sequential* task. Similarly to the *random* policy, it starts by providing the required tools, and then it repeats actions ‘a’ and ‘b’ 20 times and then cleans the workspace. Similarly to *random*, it repeats failed actions until success.

³ github.com/scazlab/human_robot_collaboration hosts the source code for the robot controllers, whereas github.com/scazlab/task-models hosts the HTM to POMDP planner.

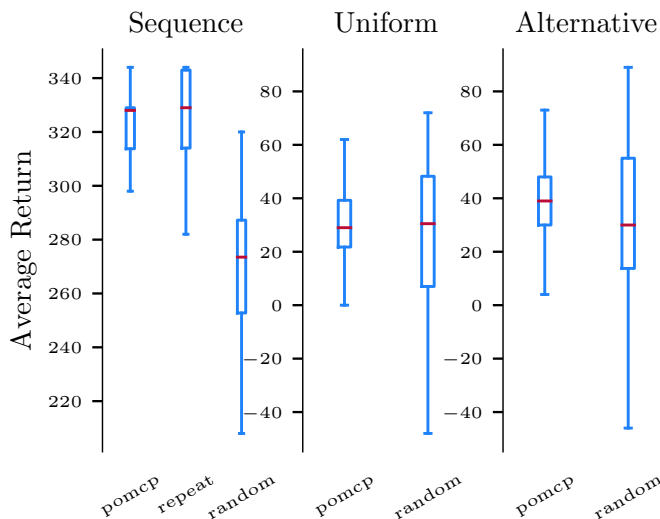


Fig. 6. The *POMCP* policy matches or outperforms ad-hoc strategies against three different HTMs (*sequential*, *uniform*, and *alternative* task). The results are the distribution of returns over 100 evaluation episodes from the two hand-coded policies (*repeat* and *random*) as well as the *POMCP* policy derived from the POMDP. On the last two tasks, the *repeat* policy fails in most of the cases and is stopped by an upper bound on the horizon of the episode. It gets a very low average return (between -300 and -400) that has been omitted from the figure to better compare the other policies.

Figure 6 presents the average return of each policy on the three conditions. Although *repeat* is very efficient on the *sequence* task, it is unfitted to the others and fails on the two other tasks. The *random* policy is suboptimal on all but the *uniform* task but can still solve them with a few failures. On the other side, the *POMCP* policy that is learned from each task model matches or outperforms the others policies. The experiment hence demonstrates that we can leverage efficient policies for each task.

2) *User Preferences*: Being able to comply with—and adapt to—custom user preferences is crucial for a robot that needs to provide the best support to its partner. A prompt and personalized response allows for a more natural interaction and a less cognitively demanding execution, which ultimately result in a more efficient collaboration. To this end, we present a system that is successfully able to account for user preferences. As detailed in Section IV-C, the participant is allowed to choose if the robot should provide support by holding during screwing or not. We compare our approach against two hard-coded policies: the most proactive strategy (i.e. always offering to hold for support) and the most conservative one (i.e. never proposing to support the human). Figure 7 demonstrates that these strategies are only optimal when they match with the expected user preference; in the intermediate scenarios, their performance degrades quickly as the uncertainty on the actual user preference increases. Conversely, our system is able to adapt to whether the human partner would like the robot to provide hold support or not, and outperforms both strategies in the majority of conditions.




Fig. 7. Comparison of the proposed strategy (red) against an ‘Always Hold’ (green) and a ‘Never Hold’ (blue) strategy during simulated interactions with varying degree of user preferences. For each of them, average return values and standard deviations with respect to the probability of the ‘Hold’ preference p_H are shown. The three strategies are tested against a single ‘Assemble Leg’ subtask (see Fig. 2b) with 20 different values of p_H , ranging from 0.0 (i.e. the human never wants the robot to hold) to 1.0 (i.e. the human always like the robot to hold). For each of the 20 different preferences, the results are averaged from 100 simulated interactions, for a total of 6000.

TABLE II
EVALUATION OF THE HUMAN-ROBOT COLLABORATION. FOR BOTH CONDITIONS, AVERAGE COMPLETION TIME AND AVERAGE NUMBER OF ROBOT ACTIONS ARE SHOWN. SEE TEXT FOR DETAILS.

	Condition A – hold	Condition B – no-hold
Average Completion Time	649 [s]	747 [s]
Avg. Number of Robot Actions	21	14

B. Live interaction with human participants

We demonstrate the proposed framework in an experimental scenario where a human participant is engaged in the joint construction of toy furniture with the Baxter robot. We define the problem detailed in Section IV-B as a sequence of eight subtasks, whose HTM is shown in Fig. 2b. All the subtasks depicted in figure require high dexterity and perception skills, and thus need to be performed by the human builder exclusively. For each of the four legs, the builder is in charge of firstly screwing the linkages (bracket and foot) onto the leg, and subsequently screwing the bracket onto the tabletop. As introduced in Section IV-B, retrieving parts and tools from their respective pool has been designed as one of the supportive actions the helper robot can choose from. Further, the Baxter robot is allowed to hold parts in order to facilitate the participant’s work, and to clean up the workstation when it deems it appropriate. For more information about the actual interaction, we refer the reader to the accompanying video, which has been summarized in Fig. 8 (full resolution available at youtu.be/OEH-DvNS0e4).

TABLE III
EXAMPLE HISTORIES OF ACTIONS AND OBSERVATIONS DURING THE INTERACTION, FOR THE TWO CONDITIONS. \hat{p}_H IS THE ESTIMATION OF THE PROBABILITY FOR THE ‘HOLD’ PREFERENCE IN THE ROBOT’S INTERNAL BELIEF.

Condition A – <i>hold</i>			Condition B – <i>no-hold</i>		
actions	obs	\hat{p}_H	actions	obs	\hat{p}_H
bring screws	none	.34	bring screws	none	.43
bring leg	none	.38	bring leg	none	.46
bring screwdriver	fail	.36	bring screwdriver	none	.44
bring top	none	.37	bring joints	fail	.38
bring joints	none	.39	bring top	none	.38
bring screwdriver	fail	.43	bring joints	none	.35
bring screwdriver	none	.42	hold	fail	.03
hold	none	.41	bring leg	none	.03
hold	none	.86	wait	none	.02
bring leg	none	.97	wait	none	.01
hold	none	.97	wait	none	.01
hold	none	.99	bring leg	none	.01
bring leg	none	1.0	wait	none	.01
hold	none	1.0	wait	none	.02
hold	none	1.0	bring leg	none	.03
bring leg	none	1.0	wait	none	.03
hold	none	1.0	wait	none	.04
hold	none	1.0	clear joints	none	.04
clear screws	none	1.0	clear screws	none	.04
clear joints	none	1.0	clear screwdriver	none	.01
clear screwdriver	none	1.0	wait	none	.01
wait	none	1.0			

The system has been evaluated with four participants; each participant performed the task in both Condition A and B, in a within-subject design (please refer Section IV-C for a description of the experimental conditions). In all the demonstrations performed, the robot was successful in providing support to the human. As shown in Table II, Condition A, when the robot was allowed to provide more support to the human—and was thus intervening more in the task, we register an overall reduction of task completion time (13.2% on average).

More interesting is to evaluate the extent to which the planner is able to recover from robot failures and policy errors. To this end, we highlight two example trajectories in Table III. They correspond to actual trajectories, one from Condition A, where the robot is expected to hold the parts, and one from Condition B, where the user will signal her preference about not wanting parts to be held. The interaction channels described in Section IV-A allow for a certain degree of flexibility for what concerns the type of communication the builder and the helper can engage into. In particular, the Baxter is allowed to gather information about the user’s preference by taking a hold action although unsure about it. This would ultimately disambiguate its uncertainty, because the builder would communicate failure in case her preference is a ‘no-hold’. For the purposes of this work, we examine here three

prototypical interactions, color-highlighted in Table III with their corresponding preference update in the third column of each table:

- *Robot-initiated failure*: as described in Section III, the robot is not allowed to directly perceive the state of the world and the progression of the task. Still, it is possible for it to detect an action failure by using its own internal sensors. If this is the case, e.g. when the robot tries to pick up the screwdriver but the gripper is empty, the system is able to re-plan its execution and re-schedule the action at a later stage. As highlighted in Table III, blue sequence, the robot does not necessarily repeat the same action right after the failure is detected, since other actions may have similar priority at that point.
- *Successful hold action* (green sequence in Table III): the robot starts with a non-zero estimation \hat{p}_H of probability for the hold preference. If, while performing an ‘hold’ action, it does not receive negative feedback from the user (observation: ‘none’), \hat{p}_H increases as it becomes more likely that the builder wants the robot to hold. This is further enforced with subsequent ‘hold’ actions.
- *User-initiated failure*: in Condition B, the system experiences an user-initiated failure while proposing to hold the part (observation: ‘fail’, red sequence in Table III). As a consequence of this, the probability of the hold action in the belief distribution \hat{p}_H decreases, and the robot will not perform this action in the future. Rather than hold parts for the user, it will wait for her to complete the action, and will move on to the next step when she communicates completion of the subtask.

One last aspect worth elaborating on is the fact that, in order for the robot to be perceived as an effective collaborator, transparency of the system during interaction is paramount. The human needs to be able to access (to a certain degree) what the robot’s internal state is, what it thinks about the task progression and, importantly, how it intends to act next. Failure to deliver transparency results in user frustration and task inefficiency. Within this context, the overlapping, redundant interaction channels (cf. Section IV-A) were beneficial in guaranteeing a transparent exchange of information between the two partners. This is particularly important in case of unexpected deviations from the robot’s nominal course of actions—that is, robot failures. Both in case of a robot-initiated failure and a human-initiated error signal (Fig. 8d), the system was able to acknowledge the user about its error state through the Baxter’s head display and/or speech utterances. In this way, it was always evident to the user that the robot failed, and eventually why it failed (in case the failure was not of robot-initiated).

VI. DISCUSSION AND FUTURE WORK

In this work, we present a system able to convert high-level hierarchical task representations into low-level robot policies. We demonstrate robustness to task representations with varying complexity, as well as a certain degree of customization with respect to task-relevant variables such as user preferences and task completion time. Further, we

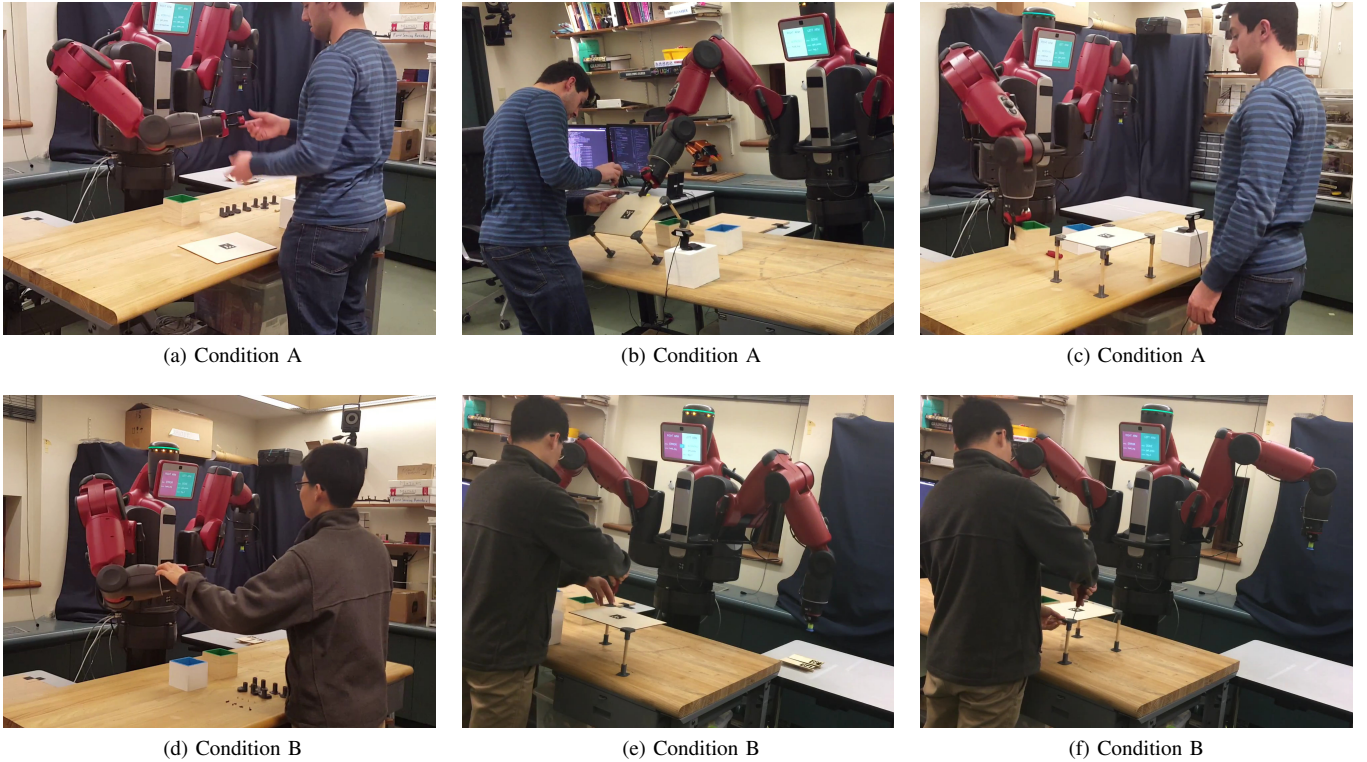


Fig. 8. Snapshots acquired during the collaborative assembly of the table in Condition A (full support, top) and Condition B (no holding required, bottom). Condition A (top): **a)** Baxter provides the tool to the participant; **b)** the robot supports the human by holding the tabletop while the human screws the leg in place; **c)** the user has finished his task and observes the robot freeing the workspace from the box of linkages. Condition B (bottom): **a)** The user signals to the robot that the holding action is not required by pressing the error button, and the robot acknowledges back by signaling into its display that it received this information; **b)** The hold action is not performed any more, but other actions such as the retrieval of the leg are still performed; **c)** The human participant completes the execution of the task without the help of the robot, as required.

introduce a novel experimental design, composed of flexible and modular constituent parts that can be easily reconfigured for a variety of different experimental scenarios. Finally, we provide demonstration of our technique in a mixed-initiative human-robot collaboration. As evident in the accompanying video, the human maintains full control throughout the task execution, but the robot acts independently, anticipates human needs, and does not wait to be told what to do. As mentioned in Section I, the paradigm in which we operate is neither to attempt implementing the ideal system that never fails and does not contemplate the occurrence of failures, nor shaping the environment in such a way that it prevents the robot to fail. Rather, we decidedly embrace the idea that robots' perception and actions are inherently faulty, and errors during operation are possible and *expected to occur*. The approach we present does not intend to compete with the optimized assembly lines that takes months to design, but provides an easy-to-deploy, reconfigurable paradigm, suitable for small and medium enterprises.

To our knowledge, this work is the first attempt at a *practical* demonstration of supportive behaviors in a realistic human-robot collaborative scenario. In addition, we fundamentally differ from past research on the topic, where collaboration typically translates to the human and the robot tasked with parallel, non overlapping subtasks and rigid, structured interactions. Rather, our experiment shows a fully integrated

interaction, where the human and the robot *physically engage in shared-environment collaboration*.

A more extensive evaluation of the scalability of the proposed framework to a broader domain of applications is the major direction for future work. Although the simulated interactions shown in Section V-A1 proved its feasibility in theory, it remains to be seen how much the approach can scale up to more complex tasks in practice. In particular, we plan on leveraging the flexibility of the HTM representation to: i) model more complex task structures; ii) apply the method to different interaction domains. Furthermore, our previous work introduced a model that allows the robot to effectively exploit basic communication capabilities in order to target the problem of task allocation and information gathering [28]. Bringing this level of interaction in the current setup is also a direction of convergence.

In this work, we assume that the high level representations our system relies on are either already available or easy enough to generate. It is however a matter of future work to explore how this assumption holds in various application domains, and whether sufficiently precise models can be learned from spoken instructions or user demonstrations. Our method also relies on existing controllers for the supportive actions. Interesting directions of future work would entail the ability of users to teach new primitives to the robot (e.g. by demonstration) and then combine these primitives into more complex task models.

One of the limitations of the live interaction experiment shown in Section V-B is that its evaluation is qualitative in nature. Nonetheless, the proposed HTM-to-POMDP method is sound, and previous work demonstrated how the same technique applied to a different POMDP allows for statistically significant results in terms of overall completion time [28]. More interestingly, the system shown here opens the door to a wide array of user studies to assess the quality and the effectiveness of the interaction between the human and the robot. The degree of proaction shown by the robot during collaboration can significantly lower the barrier to entry for non-expert users, in that users can immediately see what the robot is capable of, without over- or under-estimating its skill set. In this regard, an extensive user study would help solidifying this intuition, and assessing how useful the proposed system is in setting expectations for naive users. Finally, although our prior work showed a general user preference toward our system [28], a broader user study would prove statistical significance in terms of reduced levels of stress and cognitive load to the user. Importantly, this would also allow to highlight potential friction points that can be leveraged to better design the interaction.

REFERENCES

- [1] Google cloud speech api. <https://cloud.google.com/speech/docs/>, 2017.
- [2] Aude Billard, Sylvain Calinon, Rüdiger Dillmann, and Stefan Schaal. *Robot Programming by Demonstration*, pages 1371–1394. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [3] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.
- [4] Cynthia Breazeal. *Designing Sociable Robots*. MIT Press, 2002. ISBN 0262025108.
- [5] Cynthia Breazeal, Guy Hoffman, and Andrea Lockerd. Teaching and working with robots as a collaboration. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 3*, pages 1030–1037. IEEE Computer Society, 2004.
- [6] Andrew Garland and Neal Lesh. Learning hierarchical task models by demonstration. *Mitsubishi Electric Research Laboratory*, 2003.
- [7] S. Garrido-Jurado, R. Muñoz-Salinas, F.J. Madrid-Cuevas, and M.J. Marín-Jiménez. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition*, 47(6):2280 – 2292, 2014.
- [8] Matthew Gombolay, Ronald Wilcox, and Julie Shah. Fast Scheduling of Multi-Robot Teams with Temporospatial Constraints. In *Proceedings of Robotics: Science and Systems*, Berlin, Germany, June 2013.
- [9] Nakul Gopalan and Stefanie Tellex. Modeling and Solving Human-Robot Collaborative Tasks Using POMDPs. In *RSS Workshop on Model Learning for Human-Robot Communication*, 2015.
- [10] Jonathan Grizou, Manuel Lopes, and Pierre-Yves Oudeyer. Robot learning simultaneously a task and how to interpret human instructions. In *Development and Learning and Epigenetic Robotics (ICDL), 2013 IEEE Third Joint International Conference on*, pages 1–8. IEEE, 2013.
- [11] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- [12] Bradley Hayes and Brian Scassellati. Discovering task constraints through observation and active learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4442–4449, 2014.
- [13] Bradley Hayes and Brian Scassellati. Effective Robot Teammate Behaviors for Supporting Sequential Manipulation Tasks. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015.
- [14] Bradley Hayes and Brian Scassellati. Autonomously Constructing Hierarchical Task Networks for Planning and Human-Robot Collaboration. In *International Conference on Robotics and Automation (ICRA)*, 2016.
- [15] Pamela J. Hinds, Teresa L. Roberts, and Hank Jones. Whose Job is It Anyway? A Study of Human-robot Interaction in a Collaborative Task. *Hum.-Comput. Interact.*, 19(1):151–181, June 2004. doi: 10.1207/s15327051hci1901\&2_7.
- [16] Jesse Hoey, Pascal Poupart, Axel von Bertoldi, Tammy Craig, Craig Boutilier, and Alex Mihailidis. Automated handwashing assistance for persons with dementia using video and a partially observable Markov decision process. *Computer Vision and Image Understanding*, 114(5):503–519, 2010.
- [17] G. Hoffman and C. Breazeal. Collaboration in human-robot teams. In *In Proc. of the AIAA 1st Intelligent Systems Technical Conf.*, 2004.
- [18] G. Hoffman and C. Breazeal. Cost-Based Anticipatory Action Selection for Human Robot Fluency. *IEEE Transactions on Robotics*, 23(5):952–961, Oct 2007.
- [19] Okhtay Ilghami, Hector Munoz-Avila, Dana S Nau, and David W Aha. Learning approximate preconditions for methods in hierarchical plans. In *Proceedings of the 22nd international conference on Machine learning*, pages 337–344. ACM, 2005.
- [20] Yun Jiang, S. Moseson, and A. Saxena. Efficient grasping from RGBD images: Learning using a new rectangle representation. In *2011 IEEE International Conference on Robotics and Automation*, pages 3304–3311, May 2011.
- [21] Leslie Pack Kaelbling and Tomás Lozano-Pérez. Hierarchical task and motion planning in the now. In *Proceedings of the 1st AAAI Conference on Bridging the Gap Between Task and Motion Planning*, AAAIWS’10-01, pages 33–42. AAAI Press, 2010.
- [22] Leslie Pack Kaelbling and Tomás Lozano-Pérez. Integrated task and motion planning in belief space. *The International Journal of Robotics Research*, 32(9-10): 1194–1227, 2013.
- [23] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1):99–134, 1998.

- [24] Michael C. Koval, Nancy S. Pollard, and Siddhartha S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. *The International Journal of Robotics Research*, 35(1-3):244–264, 2016.
- [25] Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics*, 11(2):431–441, 1963. doi: 10.1137/0111030.
- [26] Stefanos Nikolaidis, Przemyslaw Lasota, Ramya Ramakrishnan, and Julie Shah. Improved human-robot team performance through cross-training, an approach inspired by human team training practices. *The International Journal of Robotics Research*, 34(14):1711–1730, 2015.
- [27] Morgan Quigley, Ken Conley, Brian P. Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y. Ng. ROS: an open-source Robot Operating System. In *ICRA Workshop on Open Source Software*, 2009.
- [28] Alessandro Roncone, Olivier Mangin, and Brian Scassellati. Transparent Role Assignment and Task Allocation in Human Robot Collaboration. *Robotics and Automation (ICRA), IEEE International Conference on*, 2017.
- [29] John Searle. *Collective Intentions and Actions*, chapter 19, pages 401–416. MIT press, 1990.
- [30] Julie Shah and Cynthia Breazeal. An empirical analysis of team coordination behaviors and action planning with application to human–robot teaming. *The Journal of the Human Factors and Ergonomics Society*, 52(2):234–245, 2010.
- [31] Julie Shah, James Wiken, Brian Williams, and Cynthia Breazeal. Improved human-robot team performance using chaski, a human-inspired plan execution system. In *Proceedings of the 6th international conference on Human-robot interaction*, pages 29–36. ACM, 2011.
- [32] G. Shani. Task-Based Decomposition of Factored POMDPs. *IEEE Transactions on Cybernetics*, 44(2): 208–216, Feb 2014.
- [33] David Silver and Joel Veness. Monte-Carlo Planning in Large POMDPs. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 2164–2172. Curran Associates, Inc., 2010.
- [34] M. Toussaint, T. Munzer, Y. Mollard, Li Yang Wu, Ngo Anh Vien, and M. Lopes. Relational activity processes for modeling concurrent cooperation. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5505–5511, May 2016.
- [35] Ronald Wilcox, Stefanos Nikolaidis, and Julie Shah. Optimization of Temporal Dynamics for Adaptive Human-Robot Interaction in Assembly Manufacturing. In *Proceedings of Robotics: Science and Systems*, Sydney, Australia, July 2012. doi: 10.15607/RSS.2012.VIII.056.



Olivier Mangin is a developmental and social robotics researcher, now a postdoctoral associate in the Social Robotics laboratory at Yale university. His interest lies in understanding how structure may appear in the sensorimotor interaction between an infant or robot and its environment. In particular he studies the development of collaborative and linguistic behaviors from observation, imitation or self exploration. Olivier completed his PhD at INRIA in Bordeaux, under the supervision of Pierre-Yves Oudeyer, after a master degree in machine learning.



focuses on the exploitation of bidirectional communication between the robot and the human in the context of human-robot collaboration and advanced manufacturing.

Alessandro Roncone is a Postdoctoral Associate at the Social Robotics Lab in Yale University. The central motivating theme of his research is endowing robots with enough communication, perception and control capabilities to enable close, natural and extended cooperation with humans. He completed his Ph.D. in Robotics, Cognition and Interaction Technologies from the Italian Institute of Technology, working on a bio-inspired peripersonal space model that improved the sensorimotor capabilities of the iCub humanoid robot. His current research



(such as autism). His other interests include humanoid robots, human-robot interaction, artificial intelligence, machine perception, and social learning.

Brian Scassellati is a Professor of Computer Science, Cognitive Science, and Mechanical Engineering at Yale University and Director of the NSF Expedition on Socially Assistive Robotics. His research focuses on building embodied computational models of human social behavior, especially the developmental progression of early social skills. Using computational modeling and socially interactive robots, his research evaluates models of how infants acquire social skills and assists in the diagnosis and quantification of disorders of social development