Robot Task Learning and Collaboration

Guy Hoffman
MIT Media Laboratory
guy@media.mit.edu

Andrea Lockerd
MIT Media Laboratory
alockerd@media.mit.edu

Abstract

In this paper we describe our work towards human-robot collaborative learning of tasks and the subsequent cooperative execution framework. In our work, a humanoid robot learns hierarchical tasks comprised of primitive action and smaller tasks, and is able to perform this task joinly with a human partner. Both the tutoring and the execution of the tasks are viewed as a collaborative discourse, building on natural human social skills and conventions.

Our approach is a goal-centric one, using goals at both the task and the action level to establish common ground in learning and collaboration. During the execution stage, dynamic meshing of subplans and self-assessment provide a turn-taking mechanism based on mutual commitment and support, resulting in a shared collaborative activity that is intuitive for a human participant.

1. Introduction

In the Robotic Life group at the MIT Media Lab we are working towards building social machines that are natural and intuitive for people to interact with. We would like to enable machines to take advantage of the multitude of social skills that humans exhibit when interacting with each other: communication, cooperation, social learning. This paper details our work towards supplying our robot, Leonardo, with the ability to learn tasks through interaction with a human teacher and then collaboratively execute these tasks with a human partner.

2. Approach

Our aim is to teach the robot a structurally complex task to later be performed collaboratively with a human. This is made possible through a goal oriented representation of tasks that affords the construction of joint intentions. This section provides an overview of the three main components of our system: hierarchical learning, goals, and joint intentions.

2.1. Hierarchical Task Learning

In our implementation the robot learns a representation of a new task, its constituent actions and sub-tasks, as well as goals associated with each of these. The task representation is such that it then affords the construction of shared plans and joint intentions in order to complete the task with a partner.

A number of social and expressive skills contribute to the robot's effectiveness in understanding and collaborating on a complex task with the human teacher. The tutoring of tasks exemplifies our approach to teaching as a collaborative discussion. Join attention is established both on the object level and on the task structure level. Leonardo uses subtle expression to indicate to the human tutor when he is ready to learn something new, and his performance of taught actions provides the tutor with immediate feedback about the robot's comprehension of the task. Envelope displays such as gaze aversion, eye contact and subtle nods and are used to segment a complex task learning structure in a natural way to the tutor. Natural key words such as ”next”, ”first” are used to indicate task structure and sequencing constraints [5].

2.2. Goal Driven Action

We believe that a goal-centric view is a fundamental feature of both teaching and collaboration. It has been repeatedly shown that humans interpret intentions based on goals [11, 7, 1] and that goals, not specific activities or motion trajectories, are what is most important in collaborative discourse. Goals provide a common ground for communication and interaction. This is particularly important in the collaborative setting, since the human partner is biased to use an intention-based psychology to interpret the agent's actions [6].

In the learning of a task, a goal is associated with each of the constituent actions as well as the task as a whole. Therefore, the task goal is more than just the conjunction of the goals of its actions and sub-tasks. Additionally, in executing the task, the task goal can be evaluated as a whole rather than evaluating each of its children's goals to determine if the task is done, improving efficiency. We found the goal driven approach crucial for both the tutoring of tasks and collaborating on them. Goals provide a common ground for action segmentation and understanding as well as for coordinating actions as part of a collaborative activity.

2.3. Joint Intention

In a collaborative task, a number of agents work together to solve a common problem. For this to take place, a joint course of action must emerge from the collection of the individual actions of the agents. In human collaboration, this does not reduce to just the sum of the individual actions derived from individual intentions, but rather is an interplay of actions inspired and affected by a joint or group intention.

Several models have been proposed to explain how joint intention results in individual intention and action to form a joint action. Searle [10] claims that collective intent and action cannot be formalized as a function of the individual intentions of the agents involved, but rather that the individual intentions are derived from their role in the common goal. He also stresses the importance of a social infrastructure to support collaboration. Bratman [2] breaks down Shared Cooperative Activity into mutual responsiveness, commitment to the joint activity and commitment to mutual support. He also introduces the idea of meshing subplans, which our project generalizes to dynamically meshing subplans. Cohen et al [4, 9] claim that a robust collaboration scheme in a changing environment with partial knowledge and beliefs requires communication, commitment to the joint task, commitment to mutual support, and dynamic meshing of subplans and action steps.

Our implementation tests some of the theoretical claims regarding joint intention. In the spirit of Bratman's SCA, we placed a high importance on communicating the robot's perceived state of the world and the task.

A world view centered on goals is important both for the ability to view a joint action with respect to a particular goal, and for the definition of individual intents based on sub-goals of the common intention. Our goal oriented task representation affords task collaboration between the robot and a human partner. Goals refer to both world and activity state, establishing common ground between the robot and the human. As a result, joint intention, attention and planning is naturally achieved. Throughout the collaboration, the human partner has a clear idea as to Leo's current singular intent as part of the joint intent.

3. Architecture

To achieve the abovementioned aims, we extended the C5M architecture [3] to handle hierarchical tasks, goal based decision making, task learning, and collaboration. This section details the implementation of each of these extensions.

3.1. Hierarchical Tasks

Tasks are represented in a hierarchical structure of actions and sub-tasks (recursively defined in the same fashion). Both tasks and actions are derived from the same *action tuple* data structure; this allows them to be used in a unified way in both the task learning and execution stages, and hence naturally affords the representation of hierarchical tasks. An *action tuple* encodes preconditions, an executable, and an until-condition. In addition to these, a task representation also encodes constraints among its actions. Currently we utilized only sequential constraints, but the representation is generic and others could be added in the future. The executable part of a task involves completing each of its child actions or tasks (unless higher-ranking goals have been achieved otherwise). A subtask executable is recursively expanded to its own actions.

3.2. Goals

The execution of tasks is driven by goals. The system currently distinguishes between two major types of goals: (a) goals that represent a state change in the world, and (b) goals that need to be executed regardless of their impact on the world (”Just Do It” [1] goals).

These types of goals differ in both their evaluation as preconditions as in their evaluations as until-conditions. A state-change goal *must* be evaluated before doing the activity to determine if it is needed. Also, the robot shows commitment towards the state change goal and proceeds to try to execute the action again if it failed to bring about the desired state change. Conversely, a ”Just Do It” goal will be performed regardless of any precondition, and will only be performed once.

3.3. Task Manager Module

The task manager module arbitrates the execution and learning of tasks. The task manager listens for ”Do:” commands from the speech understanding system. There are three possible scenarios around a task request from the human: a task is requested that Leo already knows, a task is requested but Leo needs to learn it first, or a task is requested to be performed in collaboration with the human. The person can say the following to initiate a task: ”Leo, do task 1”

1 The term ”Just Do It”, while popularized by a shoe manufacturing corporation, actually originates in Buddhist teaching, suggesting to keep our focus on out present action, and preventing thought from interfering with our state of direct presence.

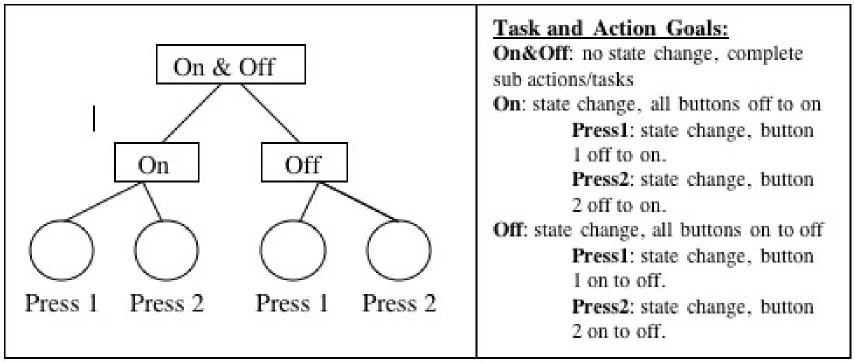

Figure 1. Hierarchical tasks and goal representation.

or ”Leo lets do task 1”, and either of these requests can also be in the form of a question (”Leo, can you do task 1?”). The task manager is responsible for recognizing these different scenarios and starting the proper execution module (also answering the person if the request was a question).

The task manager maintains a collection of known tasks. If Leo is asked to do a task on his own that he already knows, then the task manager does the execution. Execution of the task involves expanding the task’s action and subtasks onto a focus stack (in a similar way to [8]). The manager proceeds to work through the actions on the stack popping them as they are done or, upon encountering a subtask, pushing its actions to the stack.

The following two sections detail the other two scenarios, learning and collaboration. If Leo is requested to do a task in collaboration with the person, the task manager starts the collaboration module. Alternatively, if Leo is asked to do a new task, a task learning module is instantiated to learn the task.

3.4. Task Learning Module

When the task manager encounters a request for a task not in the collection of known tasks, it signals to the human that Leo needs to learn this task and proceeds to learning mode. Learning is handled recursively, such that a subtask can be learned within the learning of the larger task. While in learning mode, if the task manager receives an additional task request that is unknown, the current learning process is pushed onto a stack and an additional learning process is started. Once the subtask learning is complete, it is popped from the stack and its resulting task is added to the original learner, and the original learner resumes its process.

Task learning is best illustrated with an example; the following goes through teaching Leo to turn his buttons on and then off. The teacher begins by asking Leo to do the task, buttons on and off; to this, the robot indicates that he does not know this task and goes into learning mode. The teacher then tells him that the task starts with the task buttons on; again Leo indicates that he does not know buttons on either. As a response, the teacher begins to teach him how to do buttons on by leading him through the task. When asked to press button 1, he notices that pressing button 1 changes the state of the world, such that button 1 is now on. Hence, this is encoded as the goal of the press button 1 action, and this action is stored as part of the buttons on sub-task. The rest of the buttons on task is instructed in a similar fashion and when all of the buttons are on the teacher tells him that the buttons on task is done. At this point he notices the difference in the world before and after the buttons on task, and encodes that the goal of the buttons on task is to have all of the buttons in the ”on” state. This task is then stored as the first part of the larger buttons on and off task. The buttons off sub-task is taught in a similar way, and Leo then has a representation of the buttons off task with the goal of ending up with all of the buttons in the off state. After this he is told that the original buttons on and off task is done. Leo sees that the state of the world before and after the buttons on and off task is the same, so he assumes that the goal of the task is simply the act of doing it. The goal is encoded as a ’just-do-it’ goal. Figure 1 shows the resulting representation from the above teaching example.

Leo initiates the learning process by indicating that he does not know the requested task. Then, while in learning mode, the learning module continually records actions being performed encoding goals with these actions. When encoding the goal state of a performed action or task, Leo compares the world state before and after its execution. In the case that this action or task caused a change of state, this change is taken to be the goal. Otherwise, the goal is assumed to be ”Just Do It”. This process results in a hierarchical task representation, where a goal is encoded for each individual part of the task as well as the task as a whole. When the human indicates that the task is done, this task is added to the task managers collection of known tasks.

Since Leo shows his understanding of a newly learned

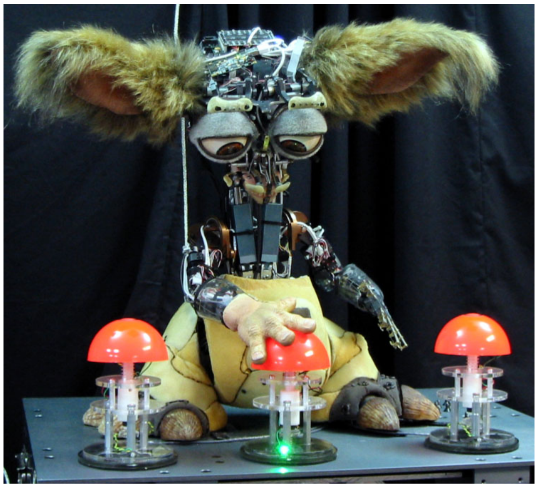

Figure 2. Leonoardo performing the steps he learns as he learns them provides the human tutor with valuable error-correcting insight in real time.

subtask or action by actually performing it (Figure 2), failure to comprehend an action or its goal is easily and naturally detected by the tutor, and we are currently working to incorporate negative feedback to correct a task representation. In a typical teacher-student interaction, errors are corrected just as they happen in the flow of the interaction; therefore, this type of error correction will be most natural for the human teacher.

3.5. Task Collaboration Module

Task collaboration is the joint execution of a common plan. When Leonardo is performing a task alone, he progresses through the task tree according to a state machine with a stack until the task's goals are achieved. When collaborating with a human partner, many new considerations come into play. In a collaborative setting, the task can (and should) be divided between the participants, the collaborator's actions need to be taken into account when deciding what to do next, mutual support is provided in cases of one participant's inability to perform a certain action, and a clear channel of communication must be used to synchronize mutual belief and maintain common ground for intentions and actions.

In our implementation we have Leonardo engage in a collaborative discourse while progressing towards achieving the joint goal. In order to make the collaboration a human-natural interaction, we have implemented many of the mechanisms that are used when humans collaborate, with a particular focus on communication, dynamic meshing of subplans, turn taking and an intuitive derivation of *I*-intentions from *We*-intentions.

3.5.1. Dynamic Meshing of Subplans We have implemented a turn taking framework in which the human collaborator and Leonardo work in together to achieve a common goal. Leo's intention system is a joint-intention model, which dynamically assigns tasks between the members of the collaboration team. The robot creates individual intentions based on his understanding of the common goal of the team, his assessment of the current task state and his understanding of his own capabilities. He is able to communicate with his human teammate about the commencement and completion of task steps and is able to recognize changes in the task environment as well as successes and failures on both his and his teammate's side. Most importantly, the robot is able to communicate to his teammate the successful completion or unattainability of a crucial task step or the complete joint action.

At every stage of the interaction, either the human should do her part in the task or Leo should do his. While usually conforming to this turn-taking approach, our system also support simultaneous action, in which the human performs an action while Leo is working on another part of the task.

To support the above capabilities, Leo derives his *I*-intentions based on a dynamic meshing of subplans according to his own actions and abilities and the actions of the human partner. If, while Leo is doing one part of the task the human completes a separate element, Leonardo will take this into account and no longer keep this on the list of things to do. Before attempting an element of the task, Leo negotiates who should complete it.

3.5.2. Self Assessment and Mutual Support Leo has the ability to evaluate his own capabilities. If he is able to complete the task element, he will offer to do so. Conversely, whenever he believes that he cannot do the action, he will ask the human for help. Since Leonardo does not, at the moment, speak, he indicates his willingness to perform an action by pointing to himself, and adopting an alert posture and facial expression (Figure 3). Analogously, when detecting an inability to perform an action assigned to him, Leo's expression indicates helplessness, as he points to the human (Figure 4). He uses gaze direction to indicate what it is he needs help with.

3.5.3. Gestural Cues A variety of gestural cues have been used in order to communicate Leo's internal state (who he thinks is doing an action, whether he believe the goal has been met) with the human. When the human partner changes the state of the world, Leo acknowledges his detecting this change by glancing shortly towards the area of change before redirecting his gaze to the human. We found this particularly valuable when the human completes part of the joint plan synchronously to Leo performing part of the

Figure 3. Leo negotiates his turn for an action he can do.

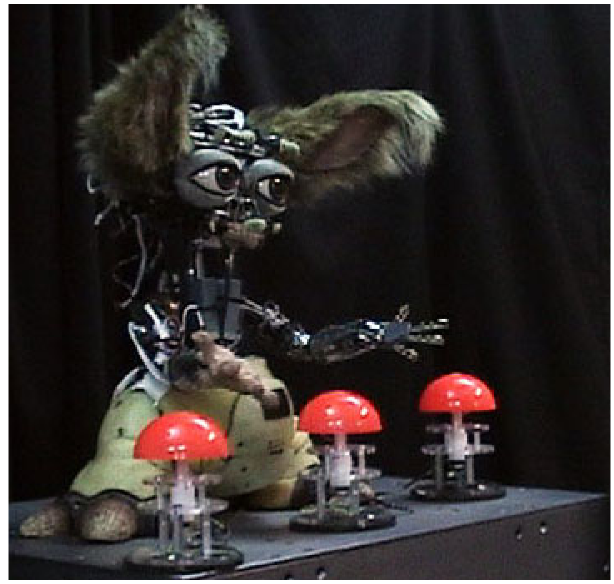

Figure 4. Leo asking for help when he encounters an action he can't perform.

task, or when the human unexpectedly changes something in the world. Leo's post-factum glance reassures the human collaborator that he has noticed the human's actions and establishes a mutual belief on the progress of the shared plan. Similary, Leo uses subtle nods while looking at his partner to indicate when he thinks a task or subtask is completed.

4. Experiments

We have conducted several trials in task learning and collaboration. Building on previously acquired speech recognition and motor skills (labeling and pressing buttons), we designed a set of tasks which involve a number of sequenced steps, such as turning a set of buttons on and then off, turning a button on as a subtask of turning all the buttons on, and turning single buttons on and off as part of "Just Do It" tasks.

In our trials, we were able to teach Leonardo the above-mentioned types of tasks, demonstrating his understanding of nested action by recalling tasks which have been learned as sub-tasks of larger activities. His understanding of task goals and the automatic classification of goals into state goals and "Just Do It" goals has been successfully demonstrated in Leo's understanding when to perform a task and how often to repeat it based on its initial success.

During training, Leo's gestural cues provided much-needed feedback enabling the tutor to realize when the robot successfully understood a task and its place in the larger scheme of things.

In the collaboration stage of our trials, the robot displayed successful meshing of subplans based on the dynamic state changes as a result of his successes, failures and the partner's actions. Leo's gestural cues provided a natural collaborative environment, informing the human partner of Leo's understanding of the task state and his attempts to take his turn. Leo's need for help displayed his understanding of his own limitations, and his use of gaze and posture served as natural cues for the human to take appropriate action in each case.

5. Conclusions

The goal of our work is to make machines more natural and rewarding for humans to interact with by incorporating natural human social skills and conventions. In this paper we have presented two important steps toward this goal: the ability to teach a task to a machine through the course of collaborative dialog and the ability to coordinate joint intentions to perform a task collaboratively. Our goal-centric approach uses goals at both the task and the action level to establish common ground in learning and collaboration.

6. Future Work

We are currently pursuing a number of extensions to the work presented here.

- As mentioned previously, in a typical teaching situation errors are corrected in the flow of the interaction as they happen with the teacher guiding the learner to the correct solution. This type of just-in-time error correction with immediate feedback from the robot will greatly improve the teaching experience for the human, and improve the accuracy of the learning as well.

- We are also working to give the system a more flexible representation of the possible goals of tasks and actions. The current scheme is rather rigid in its assumptions of what the goals of actions and tasks are. Ideally, the robot would make a number of hypothesis about what the goals could be and then generalize and become more certain about these assumptions over multiple examples.
- Currently, the task representation allows for the encoding of constraints among the constituent actions, and we are using this to specify sequential constraints. In the future we would like to specify other kinds constraints among actions, for instance representing which actions enable others.
- In the collaboration scenario, the current turn taking mechanism works by negotiating task division at each step along the way. In the future we would like to allow this negotiation to happen in advance as well, making the interaction more economical.

References

[1] D. Baldwin & J. Baird (2001), ”Discerning Intensions in Dynamic Human Action,” Trends in Cognitive Science 5(4), pp. 171-178.

[2] M. Bratman (1992). ”Shared Cooperative Activity,” The Philosophical Review , 101(2) pp. 327-341.

[3] R. Burke, D. Isla, M. Downie, Y. Ivanov, B. Blumberg (2001). ” CreatureSmarts: The Art and Architecture of a Virtual Brain,” In Proceedings of the Game Developers Conference, pp. 147-166, San Jose, CA, 2001.

[4] P. Cohen (1991). ”Teamwork,” Nous 25 , pp 487-512.

[5] P. Cohen, H. Levesque, J. Nunes, and S. Oviatt (1990), ”Task-Oriented Dialogue as a Consequence of Joint Activity,” Proceedings of PRICAI-90. Pp. 203-208.

[6] D. Dennett (1987), ”Three kinds of intentional psychology,” In The Intentional Stance. MIT Press. Chapter 3.

[7] B. Gleissner, A. Meltzoff, H. Bekkering (2000), ”Children’s coding of human action: cognitive factors influencing imitation in 3 year olds,” Developmental Science 3(4), pp. 405-414.

[8] B. Grosz & C. Sider (1990), ”Plans for Discourse,” In Cohen, Morgan, Pollack (eds) Intentions in Communication. MIT Press. Chapter 20.

[9] H. Levesque, P. Cohen, J. Nunes (1990), ”On Acting Together,” In Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90) , pp. 94–99, Boston, MA, 1990.

[10] J. Searle (1990). ”Collective Intentions and Actions,” In Cohen, Morgan and Pollack (eds.) Intentions in communication . MIT Press, Chapter 19.

[11] A. Woodward, J. Sommerville and J. Guajardo (2001), ”How infants make sense of intentional action,” In Malle, Moses, and Baldwin (eds.) Intention and Intentionality. MIT Press. Chaper 7.