

Teaching procedural flow through dialog and demonstration

Kevin Yoon, Paul E. Rybski
School of Computer Science, Carnegie Mellon University
5000 Forbes Ave., Pittsburgh, PA, 15213
{kmy,prybski}@cs.cmu.edu

Abstract—In order for robots to act as valuable assistants for non-expert users, they need to be able to learn new abilities and do so through natural methods of communication. Furthermore, it is often desirable that tasks be learned quickly without having to provide multiple demonstrations. Training should also be conducted in such a way that the user has a clear understanding of the manner in which environmental features affect the behavior of the learned activity, so that execution behavior is predictable.

We present an interactive framework for teaching a robot the flow of an activity composed of elements from a set of primitive behaviors and previously trained activities. Conditional branching and looping, order-independent activity execution, and contingency (or interrupt) actions can all be captured by our activity structures. Additional convenience functionality to aid in the training process is also provided.

By providing a natural method of communicating production rules analogous to rigid programming structures, well-defined tasks can be trained easily. We demonstrate our task training procedure on a mobile robot.

I. INTRODUCTION

In the future, robots will inevitably be employed as assistants or team partners. However, if such robots are ever to gain widespread and long term acceptance, they will need to be capable of not only learning new tasks, but also learning them from non-expert users.

We have previously introduced a method for task training via dialog and demonstration in [11]. Therein we described a collaborative natural language procedure for constructing tasks from a set of primitive behaviors and/or previously-trained tasks, which in turn could be used to build other tasks. This modular task architecture supports an expanding repertoire of abilities. Different training modes enable different features, such as the ability to attach locational context to a given command, reducing the explanatory responsibilities of the human trainer. Preconditions on task actions serve as a failure-handling mechanism that appropriately directs task flow should an action fail. The robot also engages the human in a verification dialog to resolve ambiguities in task flow and, in so doing, brings about mutual understanding of the task representation.

This understanding can be desirable, sometimes essential, in situations where the time or opportunity to provide multiple demonstrations and/or make corrections through practice trials is unavailable, and the chance of the robot exhibiting unexpected behavior due to conditions unencountered during training is unacceptable. The training dialog described herein enables the human to quickly construct rigidly-formulated

tasks where the features that affect task flow must be explicitly conveyed and not inferred. Additionally, because tasks are symbolically referenced with natural language labels, they are transferable across heterogeneous robots that share the same or a similar primitive behavior set.

In this paper, we present some enhancements and modifications to this task training technique that include the ability to capture conditional looping so that repetitive, or cyclic tasks, can be created. Interrupt events that may occur at any point during a task can also be specified to trigger contingency actions. Additionally, a new construct called a to-dolist has been added, which permits order-independent activity execution. Moreover, tasks can now be trained “on the fly” — that is, while training another task that uses it — to support a top-down design approach while still permitting the bottom-up construction of tasks. Furthermore, locational context is no longer inferred automatically as this is not always desirable in some situations. However, location-specific actions can still be specified explicitly with a simple grounding utterance.

II. RELATED WORK

Robot task learning and programming-by-demonstration (PBD) has been explored by several groups. In [1], [2], and [12], robots learn actuator trajectories or control policies from user task demonstrations. In [13], a task is built using gestures by discerning which primitive actions, from a base set of capabilities, can be combined to conduct the task demonstrated.

Our method has the ability to discern, to a limited extent, which primitive actions should be combined to execute a given task by way of inferring locational context on actions. We note, however, that this is not the main focus of our work, nor is it meant for deriving low-level control strategies. It is primarily a method by which the control flow of a task, using primitive actions and previously learned tasks, can be communicated through a training procedure employing natural interaction, thereby converging to mutual task understanding for both robot and user.

Our work is largely inspired by [8] and [7]. In [8], a mobile robot is joysticked through multiple demonstrations of a task from which it generates a generalized task representation in the form of a directed acyclic graph (DAG). The task is then pruned down to a linear sequence through teacher feedback in the form of verbal cues over multiple practice trials. In [7], a stationary humanoid robot that understands some

speech, though it is also unable to speak itself, learns tasks by communicating through gestures and facial expressions. Our approach employs a similar turn-taking framework for instruction and task refinement, but we endow the robot with the capability of speech which we believe conveys more directly the robot’s understanding of the task and guides the human more effectively in resolving ambiguities. In this way, we obviate the need to refine a learned task through practice.

Similar dialog-driven interaction mechanisms have been developed in the area of plan recognition, though primarily in the Human-Computer Interaction (HCI), as opposed to Human-Robot Interaction (HRI), domain. A plan recognition algorithm is introduced in [10] and [6] where characteristics of the collaborative setting are exploited to reduce the amount of input required of the user. This recognition strategy, however, requires some prior knowledge in the form of SharedPlans (or mutually-believed goals, actions, intentions) and a set of recipes (or action plans for achieving goals). This work differs from ours in that the goal is to help the user accomplish tasks according to perceived intent whereas we are striving to teach a robot new tasks. Our approach could potentially be used instead to build the recipes necessary for this plan recognition method to work.

In [9], an augmentation-based learning approach is described. The task structure, including conditional branching and looping, is inferred from user demonstration. Manual edits can also be made to fix incorrect task structures and constrain the induction procedure on subsequent demonstrations. Again, this approach is explored in the software application domain and there is no effort to conduct a collaborative discourse with the user for natural interaction. Additionally, in our work, branching and looping structures are explicitly and quickly communicated by the user, rather than being inferred over multiple demonstrations.

A multi-modal interface for programming tasks is described in [4] that additionally allows the user to control task priority during execution. Instruction-Based Learning [5] is similar to our work in that it uses a base set of behaviors that are associated with natural language symbolic labels and a modular architecture for symbolic tasks.

None of these works, however, describe the ability to convey branching or looping flow constructs within the task structure that are conditioned on explicitly-communicated features. Nor do they address the issue of structuring tasks for activities that need not be executed in the order in which they were communicated. This severely limits robustness and the types of tasks that can be trained. Through speech one can very compactly format instructions for execution based on detectable environmental states. No intention beliefs are maintained that may result in unexpected behavior during execution, but rather, by engaging the user in a true spoken dialog, we can quickly train tasks with clearly defined execution flow that is necessarily understood by the user.

III. SYSTEM OVERVIEW

Figure 1 depicts a simple overview of the system architecture we employ. Within the top-level behavior is an *Activity*

Selector that, upon parsing given speech commands, places the appropriate activities in the *Activity Repertoire* onto the *Current Activity List* for execution.




Fig. 1. CMAssist software architecture

An *activity* is the encompassing term for behaviors, tasks, and to-dos which are described in more detail in Section IV. The *Activity Repertoire* is the collective map that associates natural language symbolic labels to known activities. For example, “Go to” in the phrase “Go to the door” maps directly to the navigation behavior which would be put onto the *Current Activity List* with the location parameter “door”. An activity building behavior can also add new activities to the *Activity Repertoire* as will be shown in Section V-A.

Though the various activity types have differing internal structures, they are all executed by the same function form where the inputs are the *sensors* and a *command* object, and the outputs are an integer *status* flag and a new *command* object.

$$(status, command) = Activity(sensors, command)$$

The *sensors* object gives an activity module access to sensory data while *command* is an object that can be modified by an activity to store actuator commands, such as motor velocities or speech output. A single *command* object is passed through each of the activities in the *Current Activity List* so that commands requested by activities of lower priority are visible to higher priority activities. Activities can take this information into account when actuator commands need to be overridden. For example, when the obstacle avoidance behavior needs to decide whether to veer left or right to circumvent an obstacle, it can check the *command* object to see in which direction the navigation behavior was trying to drive the robot and choose to go in a similar direction. The main execution loop then involves processing all of the activities in the *Current Activity List* with the given sensory data. When the last activity on the *Current Activity List* is completed, *status* is routed back to the *Activity Selector* which determines if behaviors need to be removed from the *Current Activity List*. The *command* object is processed to drive the actuators.

The *Activity Selector* is triggered on speech input and is responsible for inserting commanded activities, removing conflicting ones, and removing completed or failed activities.

IV. ACTIVITY STRUCTURES

A. Behaviors

A behavior maps low-level sensory data to actuator trajectories in order to accomplish some high-level goal(s). The robot is assumed to be preprogrammed with some basic set of behaviors. For a mobile platform, these primitive skills might include obstacle avoidance and high-level navigation capabilities.

B. Tasks

The basic building block of a task is the task item (Figure 2). A task item consists of three main components: a (potentially empty) precondition list, an activity and a list of execution parameters, and a pointer list to subsequent task items. The precondition list contains the conditions that must be satisfied before the action can be executed. There are two types of preconditions: enabling and permanent. Enabling preconditions are evaluated only once before the task item’s activity is executed. Permanent preconditions are monitored continuously for as long as the activity is being executed. As previously mentioned, an activity can refer to a behavior, a previously-trained task, or a todoitem. Depending on the completion status of the activity (i.e. success or fail), the associated link is followed to the next task item to be executed.




Fig. 2. Task item

A task then is a temporally ordered sequence of task items captured in a directed graph structure. They can represent simple linear sequences such as in Figure 3(a). Here, the robot executes Task items 1 through N in order. Tasks can also represent conditional branching as shown in Figure 3(b). Depending on the evaluation of $\langle condition \rangle$, either Task item 2a or Task item 2b will be evaluated followed by whichever tasks follow it until the branches reconnect at Task item N. Cyclic tasks can be represented by loops as shown in 3(c). For as long as $\langle condition \rangle$ is true, Task item 2 and the subsequent task items inside the loop are executed. This is made possible by applying the while-condition as a permanent precondition on all task items inside the loop.

For some tasks it may be necessary to execute contingency activities, such as when some event occurs requiring special action and the current task be put on hold. Rather than inserting *if* and *while* statements throughout the task, the user can optionally specify contingency event-action pairs that are checked for the duration of the task execution. Unlike the previous conditional constructs, a contingency plan is not represented within the directed graph itself but is an attribute of the task structure. Each task has an associative structure that maps an interrupt event k to an action tuple (a, r) , where

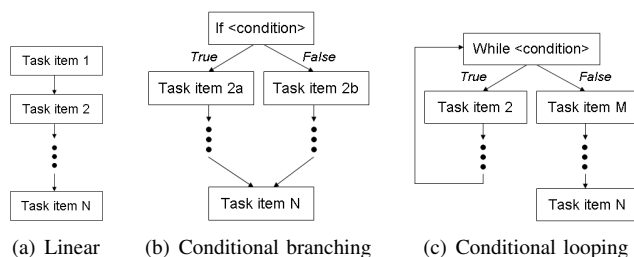


Fig. 3. Task flow structures

a is the activity to execute when k is true and r is a boolean value determining whether or not the original task should be resumed when either k is no longer true or a has completed.

C. Todolists

Todolists are a special type of activity that allows the user to specify a list of items that are to be executed in no particular order. These todoitem items, as with task items, can refer to any activity: behaviors, tasks, and other todoitems. There is nothing unique about the structure of a todoitem. It is simply a list of disconnected activities that, unlike tasks, cannot capture conditional branching and looping. It is rather the manner in which a todoitem is executed that distinguishes it from the other activities enabling it to accomplish unordered tasks as people do on a daily basis.

We currently employ a round-robin execution scheme where we iteratively loop through the list and attempt each item until it has either completed successfully or failed $maxNumTries$ times, where $maxNumTries$ is specified during training.

Clearly, some optimal scheduling strategy to minimize failed attempts could be applied here when taking into account information like estimated todoitem durations and reasons for past failures. Item priority could be an additional constraint that such a strategy might take into account. This is beyond the scope of this work where we simply provide a construct in which order-independent execution of activities is made possible.

V. TRAINING

The basic method behind the training approach we employ is allowing the user to convey production rules through the primary method of speech. Each recognized user utterance is mapped to one of three things: (1) an activity in the activity repertoire that is to be appended to the current activity structure, (2) a control structure that affects where and how subsequent activities are appended to the current activity structure, or (3) a “special” command, such as a question that the user might ask during the training procedure.

Throughout the training procedure, the robot responds with an affirmative “ok” after every user utterance to indicate understanding. The robot will also ask the user questions about parameters that were not defined when the user has finished training, thus guiding the user through dialog towards a well-defined activity structure.

A. Training Tasks

Task training is itself a behavior that can be invoked in one of two modes: dialog-only and dialog-and-observation modes. The former is invoked with the keyphrase “When I say T ” and the latter with “Let me show you what to do when I say T ”, where T is the name of the task to be trained and is typically an imperative statement. In dialog-only mode, all commands must be issued to the robot verbally. In dialog-and-observation mode, the robot invokes its following behavior such that it is always in the vicinity of the human trainer as he moves around the environment. In this manner, the robot can interpret deictic utterances like “come here”. In the previous work [11], this mode was used to automatically attach locational context to each command given by the user. In an effort to provide a framework for the training of more general tasks — where it is not necessarily appropriate to assume that actions should be executed where they were demonstrated — locational contexts are no longer assumed but can be easily and naturally anchored to subsequent commands with the “come here” phrase.

Task flow control is communicated by keyphrases summarized in Table I. An example of a user utterance that creates a conditional branching structure (Figure 3(b)) is “If you see Kevin, say ‘Hi Kevin’. Otherwise, say ‘Where is Kevin?’ before looking for Paul”. The resulting task would cause the robot to say either “Hi Kevin” or “Where is Kevin” depending on whether Kevin was detected. It would then begin the activity called *looking for Paul*.

TABLE I
TASK FLOW COMMANDS

Command	Description
“If <condition>”	Appends a conditional node to the task graph. Subsequent commands are added to <i>True</i> branch.
“Otherwise”	Causes subsequent commands to be added to the <i>False</i> branch of the current <i>if</i> node.
“before”	Connects <i>True</i> and <i>False</i> branches of current <i>if</i> node with the following command. (Ends <i>if</i> block.)
“While <condition>”	Appends a conditional node to the task graph. Subsequent commands are added to <i>True</i> branch and preconditioned on <condition>.
“After that”	Routes execution flow to the current <i>while</i> node and appends subsequent commands to the <i>False</i> branch. (Ends <i>while</i> loop.)
“Meanwhile if <condition>”	Adds a contingency event to the task object and maps it to the next activity command. If <condition> becomes true at any point during task execution, the specified activity is executed.
“Exit Task”	Appends a node that exits task with <i>success</i> flag.
“The task has failed”	Appends a node that exits task with <i>fail</i> flag.

Cyclic constructs (Figure 3(c)) can be specified with a phrase like “While Kevin is around, do a dance. After that charge batteries”. Executing the resulting task would make the robot conduct the *do a dance* activity for as long as it sees Kevin. If Kevin leaves, the loop is exited and the robot begins the *charge batteries* activity.

Contingency event-action pairs are specified with the “meanwhile” keyphrase. If we appended “Meanwhile if you see Paul, sing a song” to the previous example, the robot

would begin the *sing a song* task if Paul was detected at any time during the task (i.e. while dancing or charging batteries) and would continue to do so until the *sing a song* task completed or Paul became no longer visible. During training, the robot also asks the user if it should resume the original task after executing the contingency action.

Special utterances can be used to indicate that the task should be exited. “Exit task” and “The task has failed” create task items that when executed will terminate the task, the first with a success flag and the latter with a failure flag. (The task exits with a success flag by default even when “Exit task” is not said.) This is particularly useful when tasks are used in a to-dolist where the return status indicates whether a to-dolist item should be reattempted or not.

As can be seen, this approach to task training places more of the design burden on the user than some of the PBD techniques mentioned in Section II, but it comes with the added benefit of increased mutual task understanding between the user and robot and consequently more predictable execution behavior. Also, tasks cannot be overfit to training set conditions because task flow depends on explicitly specified features. Moreover, the natural interaction framework allows for quick and easy construction of tasks.

Figure 4 shows a simple schematic for this Task-building behavior where we can see the speech input being processed by the *Speech Parser*. Therein, we first check if the utterance is a special command, such as those shown in Table II. If it is not, then we check if it is a flow control command and add nodes or update pointers to the task under construction as appropriate. If it is not that, then we check if it corresponds to an activity that already exists in the *Activity Repertoire*. If so, then we add a task item containing the activity to the task under construction. Finally, if the user has ended the task training sequence, the robot engages the human in a verification dialog to confirm the task description by reading it back to the human and to acquire any additional information that might be necessary, such as what to do when an *if* condition does not hold and the *otherwise* case was not specified, before saving the task to the *Activity Repertoire*. The *command* object passed out of the *Speech Parser* contains speech output commands as well as any motor commands set by the Follow behavior.

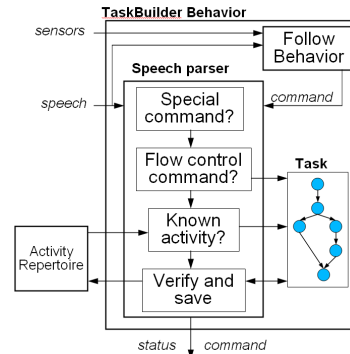


Fig. 4. Task builder

The constructs in Figure 3 can be combined and nested to create activities that can richly capture task flow. It can also be seen that activities can become arbitrarily complex. While our task training approach is well-suited for composing complex tasks from simpler subtask, the robot can provide descriptive feedback and verify with the user the flow of the trained task to minimize errors during a long and potentially confusing training sequence. In Table II are some phrases that can be understood by the robot to aid the user during the training process.

TABLE II
TRAINING HELPER FUNCTIONS

Command	Robot Function Description (T = Training mode, E = Execution mode)
“Describe T ”	T,E : Describes the task T
“What did you say?” “Can you repeat that?”	T,E : Repeats the last thing it said
“Where was I?”	T : Repeats the last two task item in the current task
Unrecognized/Misheard utterance	T : Asks if the user was referring to the name of a new task to train, and starts a new task training process if this is so. E : Robot says that it did not understand and asks the user to repeat himself.

Note that, in training mode, if the user says a phrase that is unrecognized, the robot will give the user the option of training a new task under the assumption that he may have been referring to a task that has not yet been created. In this way, the user can follow an ad-hoc, top-down approach and train tasks “on the fly” as they are required without needing to plan out all the required low-level subtasks ahead of time.

The user can say “Thank you” to simply end the training process and the learned task is saved to the *Activity Repertoire* as is. Or, by asking “Is that understood?”, the robot will dictate the task description and await confirmation from the user. If the task is correct, the robot then attempts to clarify ambiguities. Currently this involves asking the user for instructions for unspecified “otherwise” cases. If the task is incorrect, the task training procedure is restarted.

B. Training Todolists

The todolist training behavior is invoked with the keyphrase “Let’s make a todolist called L ”, where L is the name of the todolist to be trained. The user then simply lists the activities that are to be added. When finished, the user says “Thank you” to end training or asks “Is that understood?” to have the robot repeat the todolist items dictated. If the user confirms that the todolist is correct, the robot then asks the user for the number of times it should attempt to repeat failed tasks.

Unlike tasks, todolists are learned exclusively through dictation only, since todolist items themselves are typically high-level actions that can be trained as tasks.

VI. SYSTEM IMPLEMENTATION

The task training procedure was evaluated on our CMAsist¹ robot, pictured in Figure 5, that was expressly developed




Fig. 5. The CMAsist robot interacts with a user.

as our research platform for human-robot interaction. An earlier version of this task training work was demonstrated at the Robocup@Home competition in June 2006 where our team placed 2nd out of 11 teams.

The robot has modular hardware and software architectures to enable rapid prototyping and integration of new sensory, actuator, and computational components. An omnidirectional camera and stereo camera allow it to sense the presence of people wearing color-coded shirts, while the stereo camera and laser range finder together are used for navigation and obstacle avoidance. The robot can also recognize a subset of natural English language speech and speak through a Text-To-Speech (TTS) engine. These capabilities equip the robot with sufficient spatial and environmental information to execute our interactive training algorithm.

A list of relevant behaviors used by our robot is as follows:

- **Goto(x,y)/Goto(name)** Drives the robot to a location specified either by global coordinates or a location label.
- **Say(s)/Ask(s,p)** Generates speech output from the TTS engine. **Say(s)** causes the robot to speak an utterance s . **Ask(s,p)** requires that the robot identify and speak the utterance to a particular person p if present and wait for a response.
- **Follow(p)** Causes the robot to follow person p while maintaining a fixed distance of approximately 1m.
- **StateChecker(a)** Unique in that the useful output is the *status* flag rather than the *command* object, which is not modified at all. Uses sensors to calculate a *status* flag indicating whether or not an assertion a is true or false. Used by task items containing conditional statements such as *if* and *while* nodes.
- **TaskTrain(f)** Invokes the task training procedure. The **Follow** behavior is simultaneously executed if f is true causing the robot to follow the teacher and learn the task based on both the spoken utterances as well as the locations of the teacher.
- **TodolistTrain()** Invokes the todolist training procedure.

In order for the locations in the environment to be semantically meaningful as part of the training process, a map of the environment is provided to the robot which contains

¹<http://www.cs.cmu.edu/~coral/cmassist/>

semantic information in the form of location labels. For instance, the locations of named objects such as couch, table, and television can be added to the map as well as general locations of rooms such as lab or living room. This a priori information is used to ground locations that are either mentioned in the human’s speech or are visited as the human walks about the environment.

VII. EXPERIMENTAL RESULTS

The robot was trained to conduct a series of tasks that highlight the expanded capability of this task training framework. We focus on capabilities not already described in [11]. The first task is a security activity called *Patrol the lab*. This example illustrates both conditional branching and looping, makes use of an interrupt event, and demonstrates the training of tasks “on the fly”.

The transcripts for the training procedure are shown below. First, a task called *Sound the alert* (Figure 6) is trained, which is then used as the contingency action triggered when someone is detected by the robot in the *Patrol the lab* task (Figure 7).

For brevity, the “ok” feedback from the robot after every user utterance is omitted. Quoted phrases are those uttered by the user while phrases in <> are those uttered by the robot. Unquoted phrases describe what is physically happening in the scene. The numbers on the left in the transcripts are simply timestamps that denote where actions were executed on the robot’s path depicted in the corresponding scenario visualizations. Figure 8 shows the visualization for the training of *Patrol the lab*.

A. Patrol the lab

The training process for the *Patrol the lab* task is initiated in dialog-and-observation mode. Kevin says “Drive around the lab” which was not understood by the robot, so it begins a new training process. Kevin then proceeds to teach it how to drive around the lab by leading it to different locations and saying “come here” which the robot automatically converts into a *Go to X* command where *X* is the symbolic label for the current location of the human as determined from a given map. Kevin then asks “Is that understood?” to begin a verification process after which he resumes training the original *Patrol the lab* task. Having forgotten his place in the task he asks “Where was I?” and the robot reports the last two tasks items that were added. Finally, Kevin adds a contingency action, *Sound the alert*, in the event that the robot sees someone.

```

“When I say sound the alert”
“If you see Kevin”
“say hi Kevin”
“Otherwise”
“say intruder alert”
“Thank you”
<Task training is complete. Thank you.>

```

Fig. 6. Training transcript for the *Sound the alert* task

```

0 “Let me show you what to do when
  I say patrol the lab”
  <Ok, let’s begin>
6 “While Kevin is not here”
  “Drive around the lab”
  <I don’t understand. Did you want to
  teach me Drive around the lab?>
  “Yes”
  <Ok, please begin teaching me Drive around the lab>
41 Kevin walks to the door
  “Come here”
63 Kevin walks to the QRIOs
  “Come here”
81 Kevin walks to the AIBOs
  “Come here”
  “Is that understood?”
  <Yes. You said go to the door. Go to the QRIOs.
  Go to the AIBOs. Is this correct?>
  “Yes”
  <Thank you. Let’s continue training patrol the lab>
  “Where was I?”
  <You previously said while Kevin is not here,
  drive around the lab. What shall I do next?>
  “Meanwhile if you see someone”
  “sound the alert”
  <Ok. When I’m finished with that shall
  I resume patrol the lab?>
  “Yes”

```

Fig. 7. Training transcript for the *Patrol the lab* task




Fig. 8. *Patrol the lab* task training. Door, QRIOs, and AIBOs denote map locations.

By saying “Patrol the lab”, the task is invoked. Figure 9 shows the transcript of what happened and the scene is depicted in Figure 10. The robot first enters the while loop after determining that Kevin is not there - the robot cannot see people when an obstacle obstructs its view - and travels from the door, to the QRIOs, to the AIBOs, and back to the door again (Figure 10(a)). As the robot travels from the QRIOs to the AIBOs, Kevin walks in triggering the interrupt condition “meanwhile if you see someone”. The robot stops and executes *Sound the alert* causing the robot to say <Hi Kevin> (Figure 10(b)). Kevin then leaves and the robot continues with the patrol task and travels to the AIBOs. As it travels towards the door, Paul steps into view triggering the interrupt event once more. Paul is considered to be an intruder (by virtue of not being Kevin) and the robot says <Intruder alert> (Figure 10(c)). Paul then leaves and the robot again continues with the patrol task (Figure 10(d)).

```

0 "Patrol the lab"
  Goto door
13 Goto QRIOs
35 Goto AIBOs
52 Goto door
79 Goto QRIOs
103 Goto AIBOs
114 Kevin walks in. Robot stops.
    <Hi kevin>
    Kevin leaves
122 Goto door
136 Paul walks in. Robot stops.
    <Intruder alert>
    Paul leaves
152 Goto QRIOs
    Continue drive around the lab task

```

Fig. 9. Execution transcript for the *Patrol the lab* task

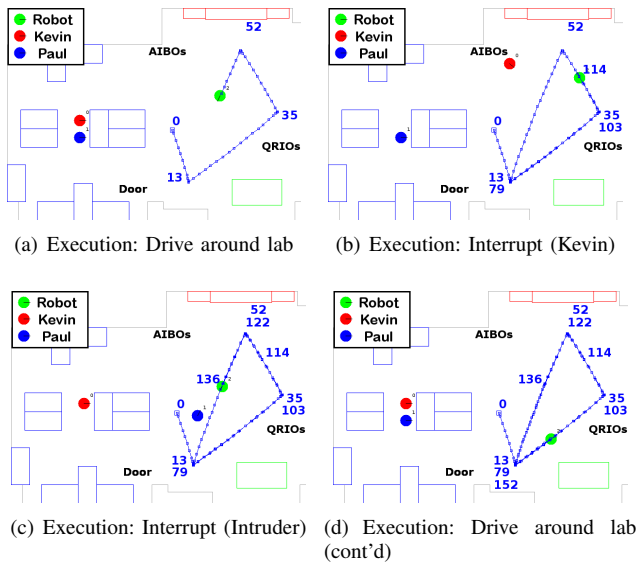


Fig. 10. *Patrol the lab* task

B. Give the lab tour

The second task is *Give the lab tour* where the robot acts as a tour guide by taking the visitors to each exhibit and then asking a human expert to talk about it in more detail. The task involves greeting visitors at the door, executing its subtasks *Give the QRIO tour* and *Give the AIBO tour*, and then saying goodbye at the center of the lab. The task is similar to the one described in [11] but includes two main differences: (1) the phrase “The task has failed” is used in the subtasks and (2) the subtasks are contained in a to-do list. This enables the robot to recognize when a task on the to-do list has failed and to reattempt it at a later time. The transcript for the training of the task *Give the QRIO tour* is shown in Figure 11. The task *Give the AIBO tour* is trained similarly except that it invokes the help of Paul, instead of Kevin, to talk about the AIBOs. The training transcript for the to-do list *Talk about the robots* and the top-level task *Give the lab tour* is shown in Figures 12 and 13, respectively. The visualization of the *Give the lab tour* is shown in Figure 14.

The execution of the *Give the lab tour* task is transcribed

```

"Let me show you what to do when
I say give the QRIO tour"
<Ok, let's begin>
Kevin goes to the QRIOs
"Come here"
"Say these are the QRIOs"
"If Kevin is there"
"ask Kevin to talk about the QRIOs"
"Otherwise"
"say unfortunately Kevin is not around to
talk about the QRIOs"
"The task has failed"
"before"
"say that concludes the QRIOs tour"

```

Fig. 11. Training transcript for the *Give the QRIO tour* task

```

"Let's make a to-do list called talk about the robots"
<Ok, let's begin>
"Give the QRIO tour"
"Give the AIBO tour"
"Is that understood?"
<Yes, you said give the QRIO tour.
Give the AIBO tour. Is this correct?>
"Yes"
<How many times shall I attempt each
item before giving up?>
"Two"
<Ok. To-do list training is complete. Thank you.>

```

Fig. 12. Training transcript for the *Talk about the robots* task

```

0 "Let me show you what to do when i
say give the lab tour"
<Ok, let's begin>
16 Kevin goes to the door
"Come here"
38 "Say welcome to the multi-robot lab"
"Talk about the robots"
50 Kevin goes to the lab center
"Come here"
"Say thank you for visiting"

```

Fig. 13. Training transcript for the *Give the lab tour* task




Fig. 14. *Give the lab tour* task training

in Figure 15 and visualized in Figure 16. After greeting the visitors at the door, the robot begins the *Talk about the robots* to-do list and goes first to the QRIOs where it finds that Kevin is not there and consequently the *Give the QRIO tour* task fails (Figure 16(a)). It then travels to the AIBOs where it sees Paul and requests that he speak about them (Figure 16(b)). When done, Paul tells the robot he is finished and the robot

```

0  "Give the lab tour"
   Goto the door
   <Welcome to the multi-robot lab>
11 Goto the QRIOs
   <These are the QRIOs. Unfortunately,
   Kevin is not around to talk about them.>
33 Goto the AIBOs
49 <These are the AIBOs. Paul, could you please
   talk about the AIBOs?>
   Paul talks about the AIBOs and then tells the
   robot that he is finished.
   Goto the QRIOs
78 <These are the QRIOs. Kevin, could
   you please talk about the QRIOs?>
   Kevin talks about the QRIOs and then
   tells the robot that he is finished.
   Goto the lab center
106 <Thank you for visiting>

```

Fig. 15. Execution transcript for the *Give the lab tour* task

decides to retry the failed items in the to-do list, i.e. *Give the QRIO tour*. This time Kevin is there and the task is completed successfully (Figure 16(c)). The tour is completed when the robot goes to the lab center and thanks the visitors for coming (Figure 16(d)).

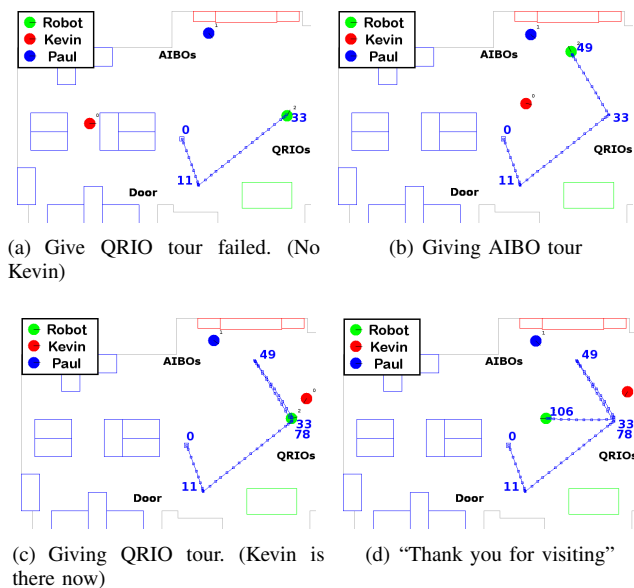


Fig. 16. *Give the lab tour* task execution

VIII. SUMMARY AND FUTURE WORK

We have presented an enhanced task training procedure that permits the user to easily communicate a rich set of task flow structures. Through dialog and observation of the user as he moves around, this framework allows for natural methods of conveying rigid production rules to construct these flow structures when training a task.

There still remains some avenues to explore in giving our robot system true utility as a personal assistant. Indeed, our work is complementary to much of the work described in

Section II and could potentially be augmented with trajectory learning and task generalization techniques described therein.

Learning of new objects, people, and locations, which was investigated to some extent in [3], would be a capability worth integrating into our system so that tasks can be conditioned on new features. This would involve improved spatial reasoning and deictic expression comprehension which would be useful in enhancing the dialog-and-observation mode of training. It would also be appropriate to symbolically parameterize tasks so that they are more generalized. The task *Give the AIBO tour* then could use the same code as the *Give the QRIO tour* task, only it would use and be conditioned on different feature parameters. This would not only decrease training times, but would require fewer resources due to code sharing.

REFERENCES

- [1] D. Bentivegna, C. Atkeson, and G. Cheng. Learning from observation and practice at the action generation level. In *IEEE International Conference on Humanoid Robots*, Karlsruhe and Munich, Germany, September/October 2003.
- [2] S. Calinon and A. Billard. Incremental learning of gestures by imitation in a humanoid robot. In *Proceedings of the 2007 ACM/IEEE International Conference on Human-Robot Interaction*, Washington, D.C., March 2007.
- [3] A. Haasch, S. Hohener, S. Huewel, M. Kleinhagenbrock, S. Lang, I. Toptsis, G.A. Fink, J. Fritsch, B. Wrede, and G. Sagerer. Biron — the bielefeld robot companion. In *Proceedings of International Workshop on Advances in Service Robotics*, pages 27–32, Stuttgart, Germany, May 2004.
- [4] S. Iba, C. J.J. Paredis, and P. K. Khosla. Interactive multi-modal robot programming. In *Proceedings of IEEE International Conference on Robotics and Automation*, Washington D.C., May 2002.
- [5] S. Lauria, G. Bugmann, T. Kyriacou, and E. Klein. Mobile robot programming using natural language. *Robotics and Autonomous Systems*, 38(3–4):171–181, 2002.
- [6] N. Lesh, C. Rich, and C. Sidner. Using plan recognition in human-computer collaboration. In *Proceedings of the Seventh International Conference on User Modelling*, Banff, Canada, June 1999.
- [7] A. Lockerd and C. Brezeal. Tutelage and socially guided robot learning. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sendai, Japan, September 2004.
- [8] M. Nicolescu and M. Mataric. Natural methods for robot task learning: Instructive demonstration, generalization and practice. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Melbourne, Australia, July 2003.
- [9] D. Oblinger, V. Castelli, and L. Bergman. Augmentation-based learning: combining observations and user edits for programming by demonstration. In *Proceedings of the International Conference on Intelligent User Interfaces*, pages 202–209, Sydney, Australia, January-February 2006.
- [10] C. Rich, C. Sidner, and N. Lesh. Collagen: Applying collaborative discourse theory to human-computer interaction. In *AI Magazine, Special Issue on Intelligent User Interfaces*, November 2001.
- [11] P. E. Rybski, K. Yoon, J. Stolarz, and M. Veloso. Interactive robot task training through dialog and demonstration. In *Proceedings of the 2007 ACM/IEEE International Conference on Human-Robot Interaction*, Washington D.C., March 2007.
- [12] J. Saunders, C. L. Nehaniv, and K. Dautenhahn. Teaching robots by moulding behavior and scaffolding the environment. In *Human-Robot Interaction*, Salt Lake City, Utah, March 2006.
- [13] R. M. Voyles, J. D. Morrow, and P. K. Khosla. Towards gesture-based programming: Shape from motion primordial learning of sensorimotor primitives. *Robotics and Autonomous Systems*, 22:361–375, November 1997.