

Task Learning Through Imitation and Human-Robot Interaction

Monica N. Nicolescu
Dept. of Computer Science and Engineering
University of Nevada, Reno
1664 W. Virginia St., MS 171
Reno, NV 89557
monica@cs.unr.edu

Maja J Matarić
Computer Science Department
University of Southern California
West 37th Place, MC 0781
Los Angeles, CA 90089-0781
mataric@cs.usc.edu

Contents

1	Task Learning Through Imitation and Human-Robot Interaction	<i>page</i> 2
1.1	Abstract	2
1.2	Introduction	2
1.3	Action-Based Representations	3
1.4	Communication by Acting - a Means for Robot-Human Interaction	5
	1.4.1 Experiments in Communication by Acting	5
	1.4.2 Discussion	7
1.5	Learning from Imitation and Additional Cues	7
	1.5.1 Experiments in Learning from Multiple Cues	9
	1.5.2 Discussion	12
1.6	Related Work	14
1.7	Conclusions	15
1.8	Acknowledgments	16
	<i>Bibliography</i>	17

1

Task Learning Through Imitation and Human-Robot Interaction

1.1 Abstract

Human skill and task teaching is a complex process that relies on multiple means of interaction and learning on the part of the teacher and of the learner. In robotics, however, task teaching has largely been addressed by using a single modality. We present a framework that uses an action-embedded representation to unify interaction and imitation in human-robot domains, thereby providing a natural means for robots to interact with and learn from humans. The representation links perception and action in a unique architecture that represents the robot's skills. The action component allows the use of implicit communication and endows the robot with the ability to convey its intentions through its actions on the environment. The perceptual component enables the robot to create a mapping between its observations and its actions and capabilities, allowing it to imitate a task learned from experiences of interacting with humans. This chapter describes a system that implements these capabilities and presents validation experiments performed with a Pioneer 2DX mobile robot learning various tasks.

1.2 Introduction

Human-robot interaction is a rapidly growing area of robotics. Environments that feature the interaction of humans and robots present a number of challenges involving robot *learning (imitative)* and *interactive* capabilities. The two problems are tightly related since, on the one hand, social interaction is often an important aspect of imitation learning and, on the other, imitative behavior enhances a robot's social and interactive capabilities. In this work we present a framework that unifies these issues, providing a natural means for robots to interact with people and to learn from interactive experiences.

We focus on two major challenges. The first is the design of robot social capabilities that allow for engagement in various types of interactions. Examples include

robot teachers [1], workers, team members [2], museum tour-guides [3], toys [4], and emotional companions [5, 6]. Designing control architectures for such robots presents various, often domain-specific, challenges.

The second challenge we address is endowing robots with the ability to learn through social interaction with humans or other robots, in order to improve their performance and expand their capabilities. Learning by imitation [7, 8, 9] provides a most natural approach to this problem; methods using gestures [10], natural language [11], and animal “clicker training” [12] have also been successfully applied.

We present an approach that unifies the above two challenges, interaction and learning in human-robot environments, by unifying perception and action in the form of *action-based interaction*. Our approach uses an architecture that is based on a set of behaviors or skills consisting of *active* and *perceptual* components. The *perceptual* component of a behavior gives the robot the capability of creating a mapping between its observations and its own actions, enabling it to learn to perform a particular task from the experiences it had while interacting with humans. The *active* component of a robot behavior allows the use of implicit communication based on action whose outcomes are invariant of the specific body performing them. A robot can thus convey its intentions by suggesting them through actions, rather than communicating them through conventional signs, sounds, gestures, or symbols with previously agreed-upon meanings. We employ these actions as a vocabulary that a robot uses to induce a human to assist it for parts of tasks that it is not able to perform on its own.

To illustrate our approach, we present experiments in which a human acts both as a collaborator and as a demonstrator for a mobile robot. The different aspects of the interaction demonstrate the robot’s learning and social abilities.

1.3 Action-Based Representations

Perception and action are the essential means of interaction with the environment. A robot’s capabilities are dependent on its available actions, and are thus an essential component of its design. The underlying control architecture we use is behavior-based [13, 14]. Behaviors are time-extended sequences of actions (e.g., go-home, avoid-obstacles) that achieve or maintain certain goals and are different than low-granularity single actions (e.g., turn-left-by-10-degrees).

Within our architecture, behaviors are built from two components, one related to perception (*Abstract behavior*), the other to action (*Primitive behavior*) (Figure 1.1). *Abstract behaviors* are explicit specifications of the activation conditions (preconditions) and effects (postconditions). *Primitive behaviors* perform the work that achieves the effects specified by those conditions. Specifically, an *abstract behavior* takes sensory information from the environment and, when its preconditions




Fig. 1.1. Structure of the inputs/outputs of an abstract and primitive behavior.

are met, activates the corresponding *primitive behavior(s)* which achieve the effects specified in its postconditions.

Using these types of behaviors, the architecture provides a natural way of representing robot tasks as hierarchical behavior networks [15] (Figure 1.2), and has the flexibility required for robust function in dynamically changing environments. This architecture is capable of computations required by more traditional symbolic architectures, but also uses behaviors continuously grounded in perception.




Fig. 1.2. Example of a behavior network

Abstract behaviors embed representations of goals in the form of abstracted environmental states. This is a key feature critical for learning from experience. To learn a task, the robot must create a mapping between its perception (observations) and its own behaviors that achieve the observed effects. This process is enabled by *abstract behaviors*, the perceptual component of a behavior, which activate each time the robot's observations match the goal(s) of a primitive behavior. This correlation enables the robot to identify its own behaviors that are relevant for the task being learned.

Primitive behaviors execute the robot's actions and achieve its goals. They are also used for communication and interaction. Acting in the environment is a form of implicit communication. By using evocative actions, people and other animals convey emotions, desires, interests, and intentions. Action-based communication has the advantage that it need not be restricted to robots or agents with a humanoid body or face: structural similarities between interacting agents are not required for successful interaction. Even if there is no direct mapping between the physical characteristics of the robot and its user, the robot can still use communication

through action to convey certain types of messages, drawing on human common sense [16].

1.4 Communication by Acting - a Means for Robot-Human Interaction

Consider a prelinguistic child who wants an out-of-reach toy. The child will try to bring a grown-up to the toy and will then point and reach, indicating his desires. Similarly, a dog will run back and forth to induce its owner to come to a place where it has found something it desires. The ability of the child and the dog to demonstrate their desires and intentions by calling a helper and mock-executing actions is an expressive and natural way to communicate a problem and need for help. The human capacity to understand these intentions is also natural and inherent. We apply the same strategy in enabling robots to communicate their desires and intentions to people.

The action-based communication approach we propose is general and can be applied on a variety tasks and physical bodies/platforms. The robot performs its task independently, but if it fails in a cognizant fashion, it searches for a human and attempts to induce him to follow it to the place where the failure occurred, and then demonstrates its intentions in hopes of obtaining help. Attracting a human to help is achieved through movement, using back-and-forth, cyclic actions. After capturing the human's attention, the robot leads the human helper to the site of the task and attempts to resume its work from the point where it failed. To communicate the nature of the problem, the robot repeatedly tries to execute the failed behavior in front of its helper. This is a general strategy that can be employed for a wide variety of failures but, notably, not for all. Executing the previously failed behavior will likely fail again, effectively expressing the robot's problem to the human observer.

1.4.1 Experiments in Communication by Acting

We implemented and tested our concepts on a Pioneer 2-DX mobile robot, equipped with two sonar rings (8 front and 8 rear), a SICK laser range-finder, a pan-tilt-zoom color camera, a gripper, and on-board computation on a PC104 stack. The robot had a behavior set that allowed it to track cylindrical colored targets (**Track (ColorOfTarget, GoalAngle, GoalDistance)**), to pick up **PickUp(ColorOfObject)**, and to drop small colored objects **Drop**. These behaviors were implemented in AYLLU [17].

In the validation experiments, we asked a person that had not worked with the robot before to be near-by during task execution and to expect to be engaged in an interaction. There is no initial assumption that people will be helpful or motivated to assist the robot. The robot is able to deal with unhelpful or misleading humans

by monitoring their presence along with its progress in the task. The following main categories of interactions emerged from the experiments:

- **uninterested:** the human was not interested, did not react to, or did not understand the robot's need for help. As a result, the robot searched for another helper.
- **interested but unhelpful:** the human was interested and followed the robot for a while, but then abandoned it. As above, the robot searched for another helper.
- **helpful:** the human was interested, followed the robot to the location of the problem, and assisted the robot. In these cases, the robot was able to finish the task.




Fig. 1.3. The human-robot interaction experiments setup

We purposefully constrained the environment used in the experiments to encourage human-robot interaction, as follows:

- *Traversing blocked gates:* the robot's task was to pass through a gate formed by two closely placed colored targets (Figure 1.3(a)), but its path was blocked by a large box. The robot expressed its intentions by executing the **Track** behavior, making its way around one of the targets. Trying to reach the desired distance and angle to the target while being hindered by box resulted in its clear manifestation of the direction it wanted to pursue, blocked by the obstacle.
- *Moving inaccessible located objects:* the robot's task was to pick up a small object which was made inaccessible by being placed in a narrow space between two large boxes (Figure 1.3(b)). The robot expressed its intentions by attempting to execute the **PickUp** behavior, lowering and opening its gripper and tilting its camera downward while approaching the object, and then moving backwards to avoid the boxes.

- *Visiting non-existing targets*: the robot's task was to visit a number of targets in a specific order (Green, Orange, Blue, Yellow, Orange, Green), in an environment where one of the targets had been removed (Figure 1.3(c)). After some time, the robot gave up searching for the missing target and sought out a human helper. The robot expressed its intentions by searching for a target, which appeared as aimless wandering. This behavior was not conducive for the human to infer the robot's goal and problem. In this and similar situations, our framework would benefit from more explicit communication.

1.4.2 Discussion

From the experimental results [18] and the interviews and report of the human subject who interacted with the robot, we derived the following conclusions about the robot's social behavior:

- **Capturing a human's attention** by approaching and then going back-and-forth is a behavior typically easily recognized and interpreted as soliciting help.
- **Getting a human to follow** by turning around and starting to go to the place where the problem occurred (after capturing the human's attention) requires multiple trials in order for the human to follow the robot the entire way. Even if interested and realizing that the robot wants something from him, the human may have trouble understanding that *following* is the desired behavior. Also, after choosing to follow the robot, if wandering in search of the place with the problem takes too long, the human gives up not knowing whether the robot still needs him.
- **Conveying intentions** by repeating a failing behavior in front of a helper is effective for tasks in which the task components requiring help are observable to the human (such as the blocked gate). However, if some part of the task is not observable (such as the missing target), the human cannot infer it from the robot's behavior and thus is not able to help (at least not without trial and error).

1.5 Learning from Imitation and Additional Cues

Learning by observation and imitation are especially effective means of human skill acquisition. As skill or task complexity increases, however, teaching typically involves increased concurrent use of multiple instructional modalities, including demonstration, verbal instruction, attentional cues, and gestures. Students/learners are typically given one or a few demonstrations of the task, followed by a set of supervised practice trials. During those, the teacher provides *feedback cues* indicating needed corrections. The teacher may also provide additional demonstrations

that could be used for *generalization*. While most of these learning and teaching tools are typically overlooked in the majority of robot teaching approaches, considering them collectively improves the imitation learning process considerably.

Toward this end, we developed a method for learning representations of high level tasks. Specifically, we augmented imitation learning by allowing the demonstrator to employ additional instructive activities (verbal commands and attentional cues) and by refining the learned representations through generalization from multiple learning experiences and through direct feedback from the teacher.

In our work, the robot is equipped with a set of skills in the form of behaviors [13, 14]; we focus on a strategy that enables it to use those behaviors to construct a high-level task representation of a novel complex, sequentially structured task. We use *learning by experienced demonstrations* in which the robot actively participates in the demonstration provided by the teacher, experiencing the task through its own sensors, an essential characteristic of our approach. We assume that the teacher knows what behaviors the robot has, and also by what means (sensors) the robot can perceive demonstrations. The advantage of *putting the robot through* the task during the demonstration is that the robot is able to adjust its behaviors (via their parameters) using the information gathered through its own sensors: the values of all behaviors' parameters are learned directly from sensory information obtained from the environment. In addition, executing the task during the demonstration provides observations that contain temporal information for proper behavior sequencing (which would otherwise be tedious to design by hand).

During the demonstration the robot follows the teacher while all of its behaviors continuously monitor the status of their postconditions (without executing any of their actions). Whenever the robot's observations match the goals of one or more primitive behavior, this means the robot has observed something it is also able to perform, and the corresponding *abstract behavior* activates, allowing the robot to learn which of its own behaviors that are relevant for the particular portion of the task being demonstrated. Feedback cues received from the teacher are used in conjunction with these observations, to eliminate any irrelevant observations.

The general idea of our task learning algorithm is to add to the robot's behavior network an instance of all behaviors whose postconditions have been detected as true during the demonstration, and during which there have been relevance signals from the teacher, in the order of their occurrence (on-line stage). At the end of the teaching experience, the intervals of time when the effects of each of the behaviors were true are known, and are used to determine if these effects were active in overlapping intervals or in sequence. Based on that information, the algorithm generates proper dependency links (*permanent, enabling or ordering*) between behaviors (off-line stage). This one-step learning process is described in detail in [18].

Through this process, the robot may learn a correct, but over-specialized version of the task. Two types of errors can occur in the learning process: *learning irrelevant steps* (false positives) and *omission of steps that are relevant* (false negatives). It is thus important that the robot can generalize over multiple demonstrations and incorporate additional *feedback* from the demonstrator. To enable a robot to *generalize* from multiple demonstrations of the same task (presented in similar or different environments), we build a task representation that encodes the specifics of each of the given examples, and also incorporates their common components [19]. However, repeated observations of irrelevant steps will also inadvertently lead the learner to include them in the learned representation. Also, limitations in robot sensing and challenging structures in the environment may prevent the robot from observing some relevant steps.

To address these issues, we allow the teacher to provide *feedback* to the robot while observing the robot’s execution of the learned task, during *practice* experiments. The teacher signals any detected errors as they occur, through appropriate feedback cues (as spoken commands). The provided feedback allows the robot to eliminate irrelevant observations and, by re-demonstrating relevant steps that were previously missed, the demonstrator enables the robot to make its learned task representation more complete [19].

1.5.1 Experiments in Learning from Multiple Cues


We implemented and tested our approach to learning with additional cues on the Pioneer 2-DX mobile robot described in Section 1.4.1. For the voice commands and feedback we used an off-the-shelf Logitech cordless headset, and the IBM ViaVoice software recognition engine. We performed two sets of robot teaching experiments to validate the key features of the proposed approach.

1.5.1.1 Generalization from a Small Number of Examples

In the first experiment set, we demonstrated the robot’s *generalization* capabilities by teaching it an object transport task in three consecutive demonstrations, performed in different environments (Figure 1.4), and designed to contain incorrect steps and inconsistencies.

The environment consisted of a set of colored cylindrical targets. The teacher lead the robot around those, instructing it when to pick up or drop a small orange box. The task to be learned was as follows: go to either the **Green (G)** or the **Light Green (LG)** targets, pick up an **Orange (O)** box, go between the **Yellow (Y)** and **Red (R)** targets, go to the **Pink (P)** target, drop the box there, then go to the **Light Orange (LO)** target, and come back to the target **Light Green**.

The shown courses of the three demonstrations illustrate that none corresponded



(a) First demonstration (b) Second demonstration (c) Third demonstration

Fig. 1.4. Structure of the environment and course of demonstration

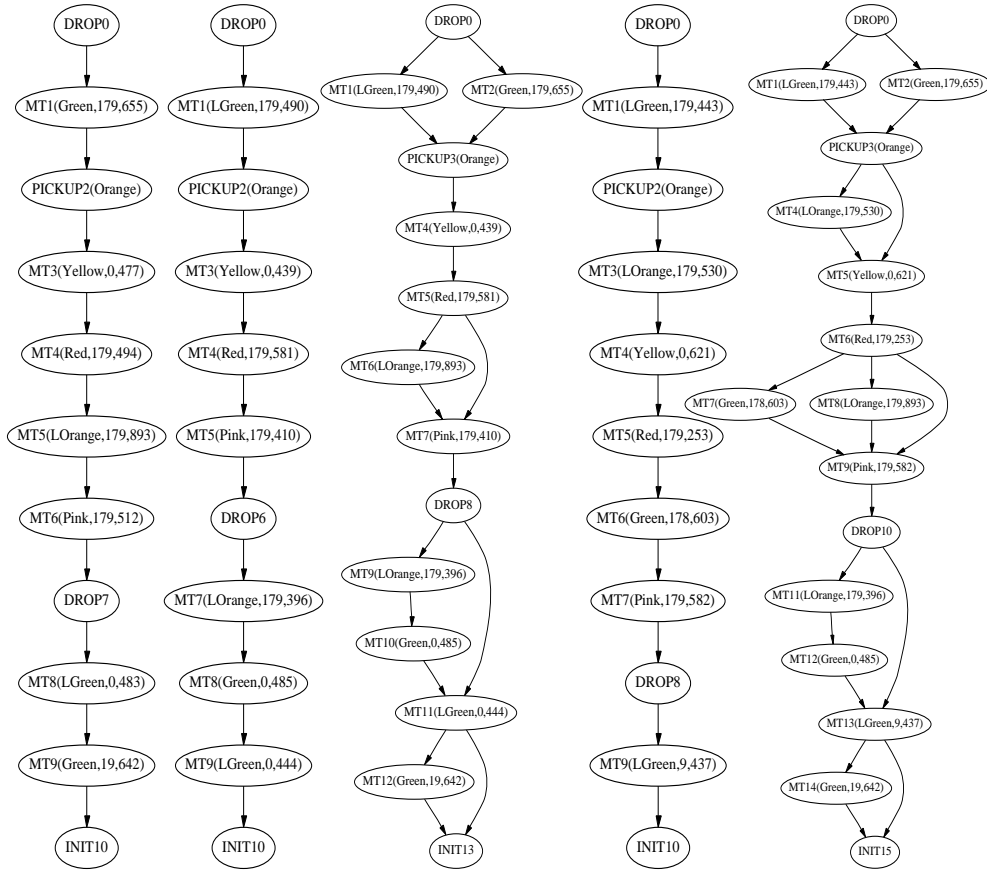
exactly to the intended task description. Some contained unnecessary steps (such as a final visit to a **Green** target in the first trial), and some had inconsistencies (such as the visits to the **Light Orange** target at various demonstration stages). Figure 1.5 shows the task representations (their topological form) obtained after each *learning* demonstration, followed by *generalization*. The topological representation of a task network is obtained by applying a *topological sort* on the behavior network graph; this representation shows the succession of behavior execution for the task. With generalization, the following types of alternative paths can be obtained:

- Both paths contain actual behaviors. For example, Figure 1.5(c) encodes the fact that both going to the **Green** or to the **Light Green** targets is acceptable for the task. Given such alternate paths, the robot chooses opportunistically, as induced by the state of the environment (e.g., go to the target seen first).
- One path is a direct link to the end of the other alternate sequence. In Figure 1.5(c), there is a direct link from **MT5(Red,...)** to **MT7(Pink,...)**, bypassing the behavior **MT6(LOrange,...)**. For such paths, the robot automatically chooses the direct path, shortcutting the alternate sequence.

The generalized representation captures the main structure of the task while correctly treating the irrelevant and inconsistent components: they are captured as parts of a bypassed alternate path that will never be executed. While irrelevant actions are thus effectively pruned, any necessary but inconsistently demonstrated steps would have to be included by different means. This is to be expected; *generalization* alone, when provided with inconsistent examples, is not sufficient for learning a correct representation. The next section shows how *practice* and *teacher feedback* can be used for solving this problem.

1.5.1.2 Learning from practice and teacher feedback

We allowed the robot to refine the previously learned task representation through practice (Figure 1.5(e)) in a different environment (Figure 1.6(a)). Figure 1.6(b) shows the robot's trajectory and the teacher's intervention (dotted). After dropping



(a) 1st demo. (b) 2nd demo. (c) Generalized task (d) 3rd demo. (e) Generalized task

Fig. 1.5. Evolution of the task representation over three successive demonstrations

the box at the destination **Pink** target, the robot started moving toward the **Light Green** target. Observing this, the teacher intervened (“**COME**”) and demonstrated to the robot the step it had missed (i.e., the visit to the **Light Orange** target). During this stage, the demonstrator also made use of informative feedback cues (“**HERE**”) to prevent the robot from considering passing by irrelevant targets (**Pink** and **Yellow** in this case) as relevant visits. The teacher signaled that it finished demonstrating the missing step (“**GO**”), after which the robot continued with and finished the task by itself. Figure 1.7(a) shows the structure of the task after this practice run. The newly added steps are marked on the graph: they also include a **Drop** behavior, as the robot had nothing in the gripper at the point of the demonstration,

and therefore the goals of this behavior were also detected as true. At execution time, the existence of this behavior had no influence, since the robot had already dropped the object.

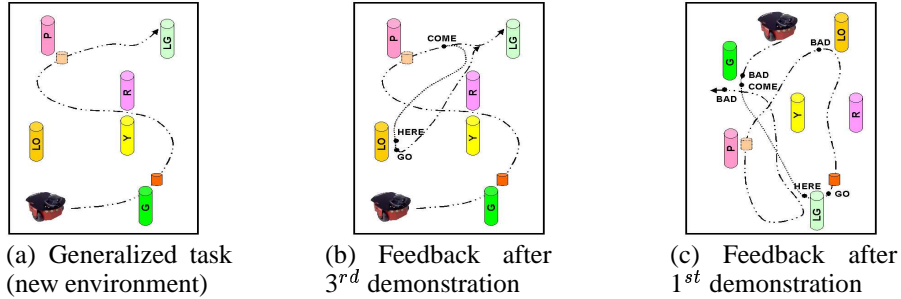


Fig. 1.6. Structure of the environment and course of task execution during practice and feedback

We now consider an alternate approach for instruction, starting from the first demonstration in the previous section (Figure 1.4(a), 1.5(a)). Assume that, for a second object transport task, the teacher considered the initial visit to a **Green** target wrong, and that the **Light Green** target should have been visited. Furthermore, the visit to the **Light Orange** target was also wrong, and not a part of the task to be learned. Figure 1.6(c) shows the trajectory of the robot and the intervention (dotted) and messages of the teacher during the robot's practice run. The effects of the teacher feedback were that: the visit to the **Green** target was replaced by a visit to the **Light Green** target, and the visits to the **Light Orange** and **Green** were removed. Figure 1.7(b) presents the structure of the task after this practice run.

The robot correctly executed the tasks learned after both of the above practice runs. In each case, it successfully adapted its task representations according to the teacher's feedback, resulting in a graph structure that matched the user's target tasks.

1.5.2 Discussion

The spectrum of tasks learnable with our method includes all tasks achievable with the robot's basic set of behaviors. If the robot is shown actions for which it does not have any representation, it will not be able to observe or learn from those experiences. For the purposes of our research, it is reasonable to accept this constraint; we are not aiming at teaching a robot new behaviors, but at showing the robot how to use its existing capabilities to perform more complicated tasks.

An additional factor influencing the performance of the proposed learning ap-

proach is the level of robotics expertise of the trainer. In prior experiments we had two non-expert users [20] serve as teachers to another robot. The results indicated that it is possible for these users to learn how to successfully teach the robot a task in a small number of trials, and with no other prior training. We have also demonstrated that the approach works for learning from robot teachers, which have significantly less expertise [21].

The *practice-feedback* experiments described above are a fast and precise method for refining previously learned tasks, since errors can be naturally indicated by the teacher immediately upon being noticed. To give the robot appropriate feedback, the teacher does not need to know the structure of the task being learned or the details of the robot's control architecture.

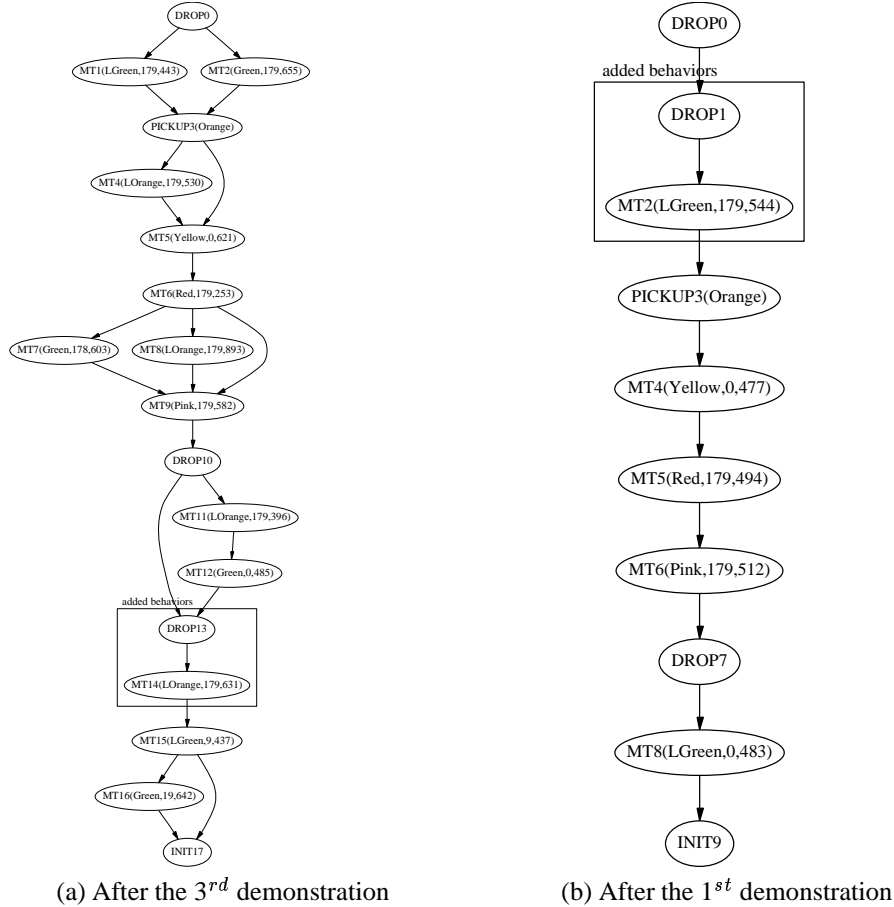


Fig. 1.7. Topologies obtained after practice and feedback

1.6 Related Work

The presented work is most related to human-robot interaction (HRI) and robot learning. Most HRI approaches so far rely on using predefined, common *vocabularies* of gestures [22], signs, or words. They can be said to be using *symbolic languages*, whose elements explicitly communicate specific meanings. One of the most important forms of implicit communication, on the other hand, is *body language*. It has been applied to humanoid robots (in particular head-eye systems), for communicating emotional states through face expressions [23] or body movements [6]. Facial expressions cannot be readily applied to mobile robots whose platforms typically have non-anthropomorphic or non-biomimetic physical structure. The work we described demonstrates how the use of implicit, action-based methods for communication and expression of intention can be applied to the mobile robot domain, despite the structural differences between mobile robots and humans, and also how communication between robots and humans can be achieved without explicit prior vocabulary sharing.

Endowing robots with the ability to imitate is of growing interest in robotics [24, 25]. In the mobile robot domain, [8] demonstrates learning to navigate a maze (learning forward, left, and right turning behaviors) by following another robot. In the humanoid robot domain, [9] used model-based reinforcement learning to speed-up learning for a system in which a 7 DOF robot arm learned the task of balancing a pole from a brief human demonstration. Other work [26, 27], related to ours in philosophy, explored imitation based on mapping observed human demonstration onto a set of behavior primitives, implemented on a 20-DOF dynamic humanoid simulation. [28] describes how the behavior primitives are learned and how they are recognized and classified in the imitation process. The key difference between that work and ours is at the level of learning: the above work focuses on learning at the level of action imitation (and thus usually results in acquiring reactive policies), while our approach enables learning of high-level, sequential tasks.

[29] demonstrate a hierarchical architecture for learning by imitation, employing HMMs to learn sequences (and thus potentially tasks) of primitive behaviors which are themselves learned and clustered. [30] describes related work on learning by imitation using two levels of primitives, the higher level employing self-organizing maps to combine (but not sequence) the lower-level primitives.

The above techniques use demonstration as the sole means for teaching and the complexity of the learned tasks is limited to reactive policies or short sequences of sensory-motor primitives. Our approach allows for learning high-level tasks that involve arbitrarily long sequences of behaviors. Methods for robot task teaching that consider additional instructive modalities in addition to demonstration have also been proposed. [31] presents an approach in which *good/not good* feedback

was given at the end of a run in which the robot performed the demonstrated skill. However, such delayed reward generates problems of credit assignment; our approach relies on immediate feedback. [32] considers fusing user intention with demonstration information as additional means for instruction. The approach enables the robot to successfully learn the correct task, but may become burdensome as the teacher is expected to provide, at each step, information on what goals he has in mind, and what actions and objects are relevant. In contrast, our approach relies solely on the teacher's observation of the robot's execution during practice.

1.7 Conclusions

We have presented an *action-based* approach to human-robot interaction and robot learning that addresses aspects of designing socially intelligent robots. The approach was shown to be effective in using implicit, action-based communication and learning by imitation to effectively interact with humans.

We argued that the means of communication and interaction for mobile robots, which do not have anthropomorphic, animal, or pet-like appearance and expressiveness, need not be limited to explicit types of interaction, such as speech or gestures. We demonstrated that simple *body language* could be used allowing the robot to successfully interact with humans and express its intentions and need for help. For a large class of intentions of the type: *I want to do "this" - but I can't*, the process of capturing a human's attention and then trying to execute the action and failing is expressive enough to effectively convey the message and obtain assistance.

Learning capabilities are essential for successful integration of robots in human-robot domains, in order to learn from human demonstrations and facilitate natural interaction with people. Due to inherent challenges of imitation learning, it is also important that robots be able to improve their capabilities by receiving additional training and feedback. Toward this end, we presented an approach that combines imitation learning with additional instructional modalities (relevant *cues*, *generalization*, *practice*), to enable a robot to learn and refine representations of complex tasks. This is made possible by the control architecture that has a perceptual component (*abstract behavior*) that creates the mapping between the observations gathered during demonstration and the robot's behaviors that achieve the same observed effects. We demonstrated these concepts on a Pioneer 2DX mobile robot, learning various tasks from demonstration, generalization, and practice.

While we believe that robots should be endowed with as many interaction modalities as is possible and efficient, this work focuses on action-based interaction as a lesser studied but powerful methodology for both learning and human-machine interaction in general.

1.8 Acknowledgments

This work was supported by DARPA Grant DABT63-99-1-0015 under the Mobile Autonomous Robot Software (MARS) program and by the ONR Defense University Research Instrumentation Program Grant N00014-00-1-0638.

Bibliography

- A. David and M. P. Ball, "The video game: a model for teacher-student collaboration," *Momentum*, vol. 17, no. 1, pp. 24–26, 1986.
- T. Matsui et al., "An office conversation mobile robot that learns by navigation and conversation," in *Proc., Real World Computing Symp.*, 1997, pp. 59–62.
- Sebastian Thrun et al., "A second generation mobile tour-guide robot," in *Proc. of IEEE, ICRA*, 1999, pp. 14–26.
- Francois Michaud and Serge Caron, "Roball, the rolling robot," *Autonomous Robots*, vol. 12, no. 2, pp. 211–222, 2002.
- Cynthia Breazeal, *Designing Sociable Robots*, MIT Press, 2002.
- Lola D. Canamero and Jakob Fredslund, "How does it feel? emotional interaction with a humanoid lego robot," in *AAAI Fall Symposium, Tech Report, FS-00-04*, Menlo Park, CA, 2000, pp. 23–28.
- Kerstin Dautenhahn and Chrystopher L. Nehaniv, Eds., *Imitation in Animals and Artifacts*, MIT Press, 2002.
- Gillian Hayes and John Demiris, "A robot controller using learning by imitation," in *Proc. of the Intl. Symp. on Intelligent Robotic Systems*, Grenoble, France, 1994, pp. 198–204.
- Stefan Schaal, "Learning from demonstration," in *Advances in Neural Information Processing Systems 9*, M.C. Mozer, M. Jordan, and T. Petsche, Eds. 1997, pp. 1040–1046, MIT Press, Cambridge.
- Richard Voyles and Pradeep Khosla, "A multi-agent system for programming robotic agents by human demonstration," in *Proc., AI and Manufacturing Research Planning Workshop*, August 1998.
- Stanislao Lauria, Guido Bugmann, Theocharis Kyriacou, and Ewan Klein, "Mobile robot programming using natural language," *Robotics and Autonomous Systems*, vol. 38, pp. 171–181, March 2002.
- Frédéric Kaplan, Pierre-Yves Oudeyer, Enikő Kubinyi, and Adam Miklúsi, "Robotic clicker training," *Robotics and Autonomous Systems*, vol. 38, pp. 197–206, March 2002.
- Maja J. Matarić, "Behavior-based control: Examples from navigation, learning, and group behavior," *Journal of Experimental and Theoretical Artificial Intelligence*, vol. 9, no. 2–3, pp. 323–336, 1997.
- Ronald C. Arkin, *Behavior-Based Robotics*, MIT Press, CA, 1998.
- Monica N. Nicolescu and Maja J. Matarić, "A hierarchical architecture for behavior-based robots," in *Proc., First Intl. Joint Conf. on Autonomous Agents and Multi-Agent Systems*, Bologna, ITALY, July 2002, pp. 227–233.

- Daniel C. Dennett, *The Intentional Stance*, MIT Press, Cambridge, 1987.
- Barry Brian Werger, "Ayllu: Distributed port-arbitrated behavior-based control," in *Proc., The 5th Intl. Symp. on Distributed Autonomous Robotic Systems*, Knoxville, TN, 2000, pp. 25–34, Springer.
- Monica N. Nicolescu and Maja J. Matarić, "Learning and interacting in human-robot domain," *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans, Special Issue on Socially Intelligent Agents - The Human in the Loop*, vol. 31, no. 5, pp. 419–430, 2001.
- Monica N. Nicolescu and Maja J. Matarić, "Natural methods for robot task learning: Instructive demonstration, generalization and practice," in *Proc., Second Intl. Joint Conf. on Autonomous Agents and Multi-Agent Systems*, Melbourne, AUSTRALIA, July 2003, pp. 241–248.
- Monica Nicolescu, *A Framework for Learning from Demonstration, Generalization and Practice in Human-Robot Domains*, Ph.D. thesis, University of Southern California, 2003.
- Monica N. Nicolescu and Maja J. Matarić, "Experience-based representation construction: Learning from human and robot teachers," in *Proc., IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, Maui, Hawaii, USA, Oct 2001, pp. 740–745.
- David Kortenkamp, Eric Huber, and R. Peter Bonasso, "Recognizing and interpreting gestures on a mobile robot," in *Proc., AAAI*, Portland, OR, 1996, pp. 915–921.
- Cynthia Breazeal and Brian Scassellati, "How to build robots that make friends and influence people," in *Proc., IROS, Kyonju, Korea*, 1999, pp. 858–863.
- Aude Billard and Kerstin Dautenhahn, "Experiments in learning by imitation - grounding and use of communication in robotic agents," *Adaptive Behavior*, vol. 7, no. 3/4, pp. 415–438, 1999.
- P. Gaussier, S. Moga, J. Banquet, and M. Quoy, "From perception-action loops to imitation processes: A bottom-up approach of learning by imitation," *Applied Artificial Intelligence Journal*, vol. 12(78), pp. 701–729, 1998.
- Maja J Matarić, "Sensory-motor primitives as a basis for imitation: Linking perception to action and biology to robotics," in *Imitation in Animals and Artifacts*, Chrystopher Nehaniv and Kerstin Dautenhahn, Eds., pp. 392–422. MIT Press, Cambridge, 2002.
- Odest Chadwicke Jenkins, Maja J Matarić, and Stefan Weber, "Primitive-based movement classification for humanoid imitation," in *Proc., First IEEE-RAS Intl. Conf. on Humanoid Robotics*, Cambridge, MA, MIT, 2000.
- Evan Drumwright, Odest C. Jenkins, and Maja J. Matarić, "Exemplar-based primitives for humanoid movement classification and control," in *IEEE Intl. Conf. on Robotics and Automation*, 2004.
- Amit Ramesh and Maja J. Matarić, "Learning movement sequences from demonstration," in *Proc. of the Intl. Conf. on Development and Learning*, MIT, Cambridge, MA, Jun 2002, pp. 203–208.
- Amit Ramesh and Maja J. Matarić, "Parametric primitives for motor representation and control," in *IEEE Intl. Conf. on Robotics and Automation*, Washington, DC, May 2002, vol. 1, pp. 863–868.
- Michael Kaiser, "Transfer of elementary skills via human-robot interaction," *Adaptive Behavior*, vol. 5, no. 3-4, pp. 249–280, 1997.
- Holger Friedrich and Rudiger Dillmann, "Robot programming based on a single demonstration and user intentions," in *Proc., 3rd European Workshop on Learning Robots*, Crete, Grece, 1995.