

Linking Perception and Action in a Control Architecture for Human-Robot Domains

Monica N. Nicolescu and Maja J Matarić
Robotics Research Laboratory
University of Southern California
941 West 37th Place, MC 0781
Los Angeles, CA 90089-0781
E-mail: monica|mataric@cs.usc.edu

Abstract

Human-robot interaction is a growing research domain; there are many approaches to robot design, depending on the particular aspects of interaction being focused on. In this paper we present an action-based framework that provides a natural means for robots to interact with humans and to learn from them. Perception and action are the essential means for a robot's interaction with the environment; for successful robot performance it is thus important to exploit this relation between a robot and its environment. Our approach links perception and actions in a unique architecture for representing a robot's skills (behaviors). We use this architecture to endow the robots with the ability to convey their intentions by acting upon their environment and also to learn to perform complex tasks from observing and experiencing a demonstration by a human teacher. We demonstrate these concepts with a Pioneer 2DX mobile robot, learning various tasks from a human and, when needed, interacting with a human to get help by conveying its intentions through actions.

1. Introduction

Human-robot interaction is an area of growing interest in Robotics. Environments that feature the interaction of humans and robots present a significant number of challenges, spawning several important research directions. These domains of human-machine co-existence form a new type of "society" in which the robot's role is essential in determining the nature of resulting interactions. In this work we focus on two major challenges of key importance for designing robots that will be effective in human-robot domains.

The first challenge we address is the design of robots that exhibit social behavior, in order to allow them to engage in

various types of interactions. This is a very large domain, with examples including teachers [5], workers, members of a team, cooperating with other robots and people to solve and perform tasks [9]. Robots can be entertainers, such as museum tour-guides [8], toys [17], pets, or emotional companions [4]. Designing control architectures for such robots presents particular challenges, in large part specific for each of these domains.

The second challenge we address is to build robots that have the ability to learn through social interaction with humans or with other robots in the environment, in order to improve their performance and expand their capabilities. Successful examples include robots imitating demonstrated tasks (such as maze learning [10] and juggling [21]) and the use natural cues (such as models of joint attention [20]) as means for social interaction.

In this paper we present an approach that unifies the two challenges above, interaction and learning in human-robot environments, by unifying *perception* and *action* in the form of *action-based interaction*. Our approach relies on an architecture that is based on a set of behaviors or skills consisting of both *active* and *perceptual* components.

The *perceptual* component of a behavior gives the robot the capability of creating a link between its observations and its own actions, which enables 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, which does not rely on a symbolic language, and instead uses actions, whose outcomes are invariant to 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 marks with previously agreed-upon meanings. We employ these actions as a vocabulary that a robot could use to induce a human to

assist it for parts of tasks that it is not able to perform on its own. The particularities of our behavior architecture are described in Section 2.

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

This paper is organized as follows. Section 2 presents the behavior representation that we are using, and the importance of the architecture for our proposed challenges. In Section 3, we present the model for human-robot interaction and the general strategy for communicating intentions, including experiments in which a robot engaged a human in interaction through actions indicative of its intentions. Section 4 describes the method for learning task representations from experienced interactions with humans and presents experimental demonstrations and validation of learning task representations from demonstration. Sections 5 and 6 discuss different related approaches and present the conclusions of the described work.

2. Behavior representation

Perception and action are the essential means of interaction with the environment. The performance and the capabilities of a robot are dependent on its available actions, and thus they are an essential component of its design. As underlying control architecture we are using a behavior-based approach [15, 1], in which time-extended actions that achieve or maintain a particular goal are grouped into behaviors, the key building blocks for intelligent, complex observable behavior. The complexity of a robot’s skills can range from elementary actions (such as “go forward”, “turn left”) to temporally-extended behaviors (such as “follow”, “go home”, etc.).




Figure 1. Structure of the inputs/outputs of an abstract and primitive behavior.

Within our architecture, behaviors are build from two components: one related to perception (*Abstract behavior*), the other to actions (*Primitive behavior*) (Figure 1). The *abstract behavior* is simply an explicit specification of the behavior’s activation conditions (i.e., preconditions), and its effects (i.e., postconditions). The behaviors that do the work

that achieves the specified effects under the given conditions are called *primitive behaviors*. An *abstract behavior* takes sensory information from the environment and, when its preconditions are met, activates the corresponding *primitive behavior(s)*, which achieve the effects specified in its postconditions.

This architecture provides a simple and natural way of representing robot tasks in the form of behavior networks [19], and also has the flexibility required for robust function in dynamically changing environments. Figure 2 shows a generic behavior network.




Figure 2. Example of a behavior network

The *abstract behaviors* embed representations of a behavior’s goals in the form of abstracted environmental states. This is a key feature of our architecture, and a critical aspect for learning from experience. *In order to learn a task the robot has to create a link between perception (observations) and the actions that would achieve the same observed effects.* This process is enabled by the *abstract behaviors*, the perceptual component of a behavior. This component fires each time the observations match a primitive’s goals, allowing the robot to identify during its experience the behaviors that are relevant for the task being learned.

The *primitive behaviors* are the active component of a behavior, executing the robot’s actions and achieving its goals. Acting in the environment is a form of implicit communication that plays a key role in human interaction. 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 the interacting agents are not required to achieve successful interaction. Even if there is no exact mapping between a mobile robot’s physical characteristics and those of a human user, the robot may still be able to convey a message, since communication through action also draws on human common sense [6]. In the next section we describe how our approach achieves this type of communication.

3. Communication by acting - a means for robot-human interaction

Our goal is to develop a model of interaction with humans that would allow a robot to induce a human to assist it

by being able to express its intentions in a way that humans could easily understand. We first present a general example that illustrates the basic idea of our approach.

Consider a prelinguistic child who wants a toy that is out of his reach. To get it, the child will try to bring a grown-up to the toy and will then point and even try to reach it, indicating his intentions. 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 intentions by calling a helper and mock-executing an action is an expressive and natural way to communicate a problem and need for help. The capacity of a human observer to understand these intentions from exhibited behavior is also natural since the actions carry intentional meanings, and thus are easy to understand.

We apply the same strategy in the robot domain. The action-based communication approach we propose for the purpose of suggesting intentions is general and can be applied across different tasks and physical bodies/platforms. In our approach, a 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 demonstrates its intentions in hopes of obtaining help. Next, we describe how this communication is achieved.

Immediately after a failure, the robot saves the current state of the task execution (failure context), in order to be able to later restart execution from that point.



Figure 3. Behavior network for calling a human

Next, the robot starts the process of finding and luring a human to help. This is implemented as a behavior-based system, which uses two instances of a **Track(Human, angle, distance)** behavior, with different values of the *Distance* parameter: one for getting close (50cm) and one for getting farther (1m) (Figure 3). As part of the first tracking behavior, the robot searches for and follows a human until he stops and the robot gets sufficiently close. At that point, the preconditions for the second tracking behavior are active, so the robot backs up in order to get to the farther distance. Once the outcomes of this behavior have been achieved (and detected by the **Init** behavior), the robot re-instantiates the network, resulting in a back and forth *cycling* behavior, much like a dog’s behavior for enticing a human to follow it. When the detected distance between the robot and the human becomes smaller than the values of the *Distance* parameter for any one of its **Track** behaviors for some period of time, the cycling behavior is terminated.

The **Track** behavior enables the robot to follow colored targets at any distance in the [30, 200] cm range and any angle in the [0, 180] degree range, by merging the information from the camera and the laser range-finder. Thus, the behavior gives the robot the ability to keep track of positions of objects around it, even if they are not currently visible, akin to working memory.

After capturing the human’s attention, the robot switches back to the task it was performing, from the point where it failed, while making sure that the human is following. This is achieved by adjusting the speed of the robot such that the human follower is kept within desirable range behind the robot. If the follower is lost, the robot starts searching again for another helper. After a few experiences with *unhelpful* humans, the robot will again attempt to perform the task on its own. If a human provides useful assistance, and the robot is able to execute the previously failed behavior, the robot continues with task execution as normal.

Thus, the robot retries to execute its task from the point where it has failed, while making sure that the human helper is near by. Executing the previously failed behavior will likely fail again, effectively expressing to the human the robot’s problem.

In the next section we describe the experiments we performed to test the above approach to human-robot interaction, involving cases in which the human is helpful, unhelpful, or uninterested.

3.1. Experiments on Robot Interacting with humans - Communication by Acting

The experiments that we present in this section focus on performing actions as a means of communicating intentions and needs. Initially, the robot (which has a typical mobile robot form entirely different from that of the human) was given a behavior set that allowed it to track colored targets, open doors, pick up, drop, and push objects. The behaviors were implemented using AYLLU [22], an extension of the C language for development of distributed control systems for mobile robot teams.

We tested our concepts on a Pioneer 2-DX mobile robot, equipped with two rings of sonars (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.

In order to test the interaction model we described above, we designed a set of experiments in which the environment was changed so that the robot’s execution of the task became impossible without some outside assistance.

The failure to perform any one of the steps of the task induced the robot to seek help and to perform evocative actions in order to catch the attention of a human and get him to the place where the problem occurred. In order to communicate the nature of the problem, the robot repeatedly

tried 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. However, as demonstrated in our third example below, there are situations for which this approach is not sufficient for conveying the message about the robot's intent. In those, explicit communication, such as natural language, is more effective. We discuss how different types of failures require different modes of communication for help.

In our validation experiments, we asked a person that had not worked with the robot before to be close during the tasks execution and expect to be engaged in interaction. During the experiment set, we encountered different situations, corresponding to different reactions of the human in response to the robot. We can group these cases into the following main categories:

- **uninterested:** the human was not interested in, did not react to, or did not understand the robot's calling for help. As a result, the robot started to search for another helper.
- **interested, unhelpful:** the human was interested and followed the robot for a while but then abandoned it. As in the previous case, when the robot detected that the helper was lost, it started to look for another one.
- **helpful:** the human followed the robot to the location of the problem and assisted the robot. In these cases the robot was able to finish the execution of the task, benefiting from the help it had received.


We purposefully constrained the environment in which the task was to be performed, in order to encourage human-robot interaction. The helper's behavior, consequently, had a decisive impact on the robot's task performance: when uninterested or unhelpful, failure ensued either due to exceeding time constraints or to the robot giving up the task after trying for too many times. However, there were also cases when the robot failed to find or entice the human to come along, due to visual sensing limitations or the robot failing to expressively execute its *calling* behavior. The few cases in which a failure occurred despite the assistance of a helpful human are presented below, along with a description of each of the three experimental tasks and overall results.

3.1.1. Traversing blocked gates

In this section we discuss an experiment in which a robot is given a task of traversing gates formed by two closely placed colored targets (see Figure 4(a)). The environment is arranged such that the path between the targets is blocked by a large box that prevents the robot from going through.

Expressing intentionality of performing this task is done by executing the **Track** behavior, which allows the robot to make its way around one of the targets. While trying to reach the desired distance and angle to the target, hindered

by the large box, the robot shows the direction it wants to go in, which is blocked by the obstacle.



(a) Going through a gate (b) Picking up an inaccessible box (c) Visiting a missing target

Figure 4. The human-robot interaction experiments setup

We performed 12 experiments in which the human proved to be helpful. Failures in accomplishing the task occurred in three of the cases, in which the robot could not get through the gate even after the human had cleared the box from its way. For the rest of the cases the robot successfully finished the task with the human's assistance.

3.1.2. Moving inaccessible located objects

The experiment described in this section involves moving objects around. The robot is supposed to pick up a small object, close to a big blue target. In order to induce the robot to seek help, we placed the desired object in a narrow space between two large boxes, thus making it inaccessible to the robot (see Figure 4(b)).

The robot expresses the intentions of getting the object by simply attempting to execute the corresponding **PickUp** behavior. This forces the robot to lower and open its gripper and tilt its camera down when approaching the object. The drive to pick up the object is combined with the effect of avoiding large boxes, causing the robot to go back and forth in front of the narrow space and thus convey an expressive message about its intentions and its problem.

From 12 experiments in which the human proved to be helpful, we recorded two failures in achieving the task. These failures were due to the robot losing track of the object during the human's intervention and being unable to find it again before the allocated time expired. For the rest of the cases the help received allowed the robot to successfully finish the task execution.

3.1.3. Visiting non-existing targets

In this section we present an experiment that does not fall into the category of the tasks mentioned above and is an ex-

ample for which the framework of communicating through actions should be extended to include more explicit means of communication. Consider a task of visiting a number of targets, in a given order (Green, Orange, Blue, Yellow, Orange, Green), in which one of the targets has been removed from the environment (Figure 4(c)).

The robot gives up after some time of searching for the missing target and goes to the human for help. By applying the same strategy of executing in front of the helper the behavior that failed, the result will be a continuous wandering in search of the target from which it is hard to infer what the robot’s goal and problem are. It is evident that the robot is looking for something - but without the ability to name the missing object, the human cannot intervene in a helpful way.

3.2. Discussion

The experiments presented above demonstrate that implicit yet expressive action-based communication can be successfully used even in the domain of mobile robotics, where the robots cannot utilize physical structure similarities between themselves and the people they are interacting with. However, as our third experiment showed, there are situations in which actions alone are not sufficient for conveying the robot’s intent. This is due to the fact that the failure the robot encountered has aspects that could not be expressed by only repeating the unsuccessful actions. For those cases we should employ explicit forms of communication, such as natural language, to convey the necessary information.

From the results, our observations, and the report of the human subject interacting with the robot throughout the experiments, we derive the following conclusions about the various aspects of the robot’s social behavior:

- **Capturing a human’s attention** by approaching and then going back and forth in front of him 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 completely follow the robot the entire way. This is due to several reasons: first, even if interested and realizing that the robot wants something from him, the human may not actually believe that he is being called by a robot in a way in which a dog would do it and does not expect that *following* is what he should do. Second, after choosing to go with the robot, if wandering in search of the place with the problem takes too much time, the human gives up not knowing whether the robot still needs him.
- **Conveying intentions** by repeating the actions of a failing behavior in front of a helper is easily achieved for tasks

in which all the elements of the behavior execution are observable to the human. Upon reaching the place of the robot’s problem, the helper is already engaged in interaction and is expecting to be shown something. Therefore, seeing the robot trying and failing to perform certain actions is a clear indication of the robot’s intentions and need for assistance.

4. Learning from human demonstrations

Automating the robot controller design process is a topic of particular interest for robotic researchers; its goal is to allow both specialized and non-specialized users to easily “program” the robots according to their needs. A natural approach to this problem is the use of teaching by demonstration. Instead of having to write, by hand, a controller that achieves a particular task, we allow a robot to automatically build it from observation or from the experience it had while interacting with a teacher. It is the latter approach that we will consider in this work, as a means for transfer of task knowledge from teachers to robots.

We assume that the robot is equipped with a set of behaviors, also called *primitives*, which can be combined into a variety of tasks. We then focus on a learning strategy that would help a robot build high-level task representation that will achieve the goals demonstrated by a teacher through the activation of the existing behavior set. We do not attempt to reproduce exact trajectories or actions of the teacher, but rather learn the task in terms of its high-level goals.

In our particular approach to learning, we use *learning by experienced demonstrations*. This implies that the robot actively participates in the demonstration provided by the teacher, by following the human, and experiencing the task through its own sensors. Thus, our approach is once again action-based: the robot has to perform the task in order to learn it. This is an essential characteristic of our approach, and is what is providing the robot the data necessary for learning. In the mobile robot domain the experienced demonstrations are achieved by following of and interacting with the teacher. The advantage of “putting the robot through” the task during the demonstration is that the robot is able to adjust its behaviors (through their parameters) using the information gathered through its own sensors. In contrast, if the task were designed by hand, a user would have to determine those parameter values. Furthermore, if the robot were merely observing but not executing the task, it would also have to estimate the parameter values at least for the initial trial or set of trials. In addition to experiencing parameter values directly, the execution of the behaviors provides observations that contain temporal information for proper behavior sequencing, which would be tedious to design by hand for tasks with long temporal sequences.

An important challenge for a learning method that is

based on robot’s observations is to distinguish between the relevant and irrelevant information that the robot is perceiving. In our architecture, the *abstract behaviors* help the robots significantly in pruning the observations that are not related to their own skills, but it is still impossible to determine exactly what is really relevant for a particular task. For example, while teaching a robot to go and pick up the mail, a robot can detect numerous other aspects along its path (e.g., passing a chair, meeting another robot, etc.). However, these observations should not be included in the robot’s learned task, as they are irrelevant for getting the mail.

To have a robot learn a task correctly in such conditions, the teacher needs a means of providing the robot with additional information than just the demonstration experience. In our approach, the teacher is allowed to signal through gestures (by showing a colored marker) the moments in time when the environment presents aspects relevant to the task. While this allows the robot to distinguish some of the irrelevant observations, it still may not help it to perfectly learn the task. For this, methods such as multiple demonstrations and generalization techniques can be applied. We are currently investigating these methods as a future extension to this work.

The general idea of the algorithm is to add to the network task representation an instance of all behaviors whose postconditions have been true during the demonstration, and during which there have been signals from the teacher, in the order of their occurrence. At the end of the teaching experience, the intervals of time when the effects of each of the behaviors have been true are known, and are used to determine if these effects have been active in overlapping intervals or in sequence. Based on the above information, the algorithm generates the proper network links (i.e., precondition-postcondition dependencies). This learning process is described in more detailed in [18].

4.1. Experimental results - learning in clean environments


We performed three different experiments in a 4m x 6m arena, in which only the objects relevant to the tasks were present. During the demonstration phase, a human teacher led the robot through the environment while the robot recorded its observations relative to the postconditions of its behaviors.

We repeated these *teaching* experiments more than five times for each of the demonstrated tasks, to validate that our learning algorithm reliably constructs the same task representation for the same demonstrated task. Next, using the behavior networks constructed during the robot’s observations, we performed experiments in which the robot reliably repeated the task it had been shown and had learned. We

tested the robot in executing the task five times in the same environment as the one in the learning phase, and also five times in a changed environment. We present the details and the results for each of the tasks in the following sections.

4.1.1. Learning to visit targets in a particular order

The goal of this experiment was to teach the robot to reach a set of targets in the order indicated by the arrows in Figure 5(a). The robot’s behavior set contains a **Tracking** behavior, parameterizable in terms of the colors of targets that are known to the robot. Therefore, during the demonstration phase, different instances of the same behavior produced output according to their settings.



(a) Experimental setup (1) (b) Experimental setup (2) (c) Approximate robot trajectory

Figure 5. Experimental setup for the “Visit targets” task




Figure 6. Task representation learned from the demonstration of the “Visit targets” task

Figure 6 shows the behavior network the robot constructed as a result of the above demonstration.

More than five trials of the same demonstration were performed in order to verify the reliability of the network generation mechanism. All of the produced controllers were identical and validated that the robot learned the correct representation for this task.

4.1.2. Learning to slalom

In this experiment, the goal was to teach a robot to slalom through four targets placed in a line, as shown in

Figure 7(a). We changed the size of the arena to 2m x 6m for this task.




Figure 7. The “Slalom” task: (a) Experimental setup; (b) Approximate robot trajectory; (c) Task representation learned from the demonstration of the Slalom task

During 8 different trials the robot learned the correct task representation as shown in the behavior network from Figure 7(c).

We performed 20 experiments, in which the robot correctly executed the slalom task in 85% of the cases. The failures consisted of two types: 1) the robot, after passing one “gate,” could not find the next one due to the limitations of its vision system; and 2) the robot, while searching for a gate, turned back toward the already visited gates. Figure 7(b) shows the approximate trajectory of the robot successfully executing the slalom task on its own.

4.1.3. Learning to traverse “gates” and move objects from one place to another

The goal of this experiment was to extend the complexity of the task to be learned by adding to it object manipulation. For this, the robot used its behaviors for picking up and dropping objects in addition to the behaviors for navigating and tracking, already described.

The setup for this experiment is presented in Figure 8(a). Note the small orange box close to the green target. In order to teach the robot that the task is to pick up the orange box placed near the green target (the source), the human led the robot to the box, and when sufficiently near it, placed the box between the robot’s grippers. After leading the robot through the “gate” formed by the blue and yellow targets, when reaching the orange target (the destination), the human took the box from the robot’s gripper. The learned behavior network representation is shown in Figure 9. Since the robot started the demonstration with nothing in the gripper, the effects of the **Drop** behavior were met, and thus an instance of that behavior was added to the network. This ensures correct execution for the case when the robot might

start the task while holding something: the first step would be to drop the object being carried.




Figure 8. The “Object manipulation” task: (a) Traversing gates and moving objects; (b) Approximate trajectory of the robot




Figure 9. Task representation learned from the demonstration of the “Object manipulation” task

The ability to track targets within a [0, 180] degree range allows the robot to learn to naturally execute the part of the task involving going through a gate. This experience is mapped onto the robot’s representation as follows: “track the yellow target until it is at 180 degrees (and 50cm) with respect to you, then track the blue target until it is at 0 degrees (and 40cm).” At execution time, since the robot is able to track both targets even after they disappeared from its visual field, the goals of the above **Track** behaviors were achieved with a smooth, natural trajectory of the robot passing through the gate.

Due to the increased complexity of the task demonstration, in 10% of the cases (out of more than 10 trials) the behavior network representations built by the robot were not completely accurate. The errors represented special-

ized versions of the correct representation, such as: **Track** the green target from a certain angle and distance, followed by the same **Track** behavior but with different parameters - only the last was in fact relevant.

The robot correctly executed the task in 90% of the cases. The failures were all of the type involving exceeding the allocated amount of time for the task. This happened when the robot failed to pick up the box because it was too close to it and thus ended up pushing it without being able to perceive it. This failure results from the undesirable arrangement and range of the robot’s sensors, not to any algorithmic issues.

4.1.4. Discussion

The results obtained from the above experiments demonstrate the effectiveness of using human demonstration combined with our behavior architecture as a mechanism for learning task representations. The approach we presented allows a robot to automatically construct such representations from a single demonstration. The summary of the experimental results is presented in Table 1. Furthermore, the tasks the robot is able to learn can embed arbitrarily long sequences of behaviors, which are encoded within the behavior network representation.

Table 1. Summary of the experimental results.

Experiment name	Trials	Successes	
		Nr.	Percent
Six targets (learning)	5	5	100 %
Six targets (execution)	5	5	100 %
Slalom (learning)	8	8	100 %
Slalom (execution)	20	17	85 %
Object move (learning)	10	9	90 %
Object move (execution)	10	9	90 %

Analyzing the task representations the robot built during the experiments above, we observe the tendency toward over-specialization. The behavior networks the robot learned enforce that the execution go through all demonstrated steps of the task, even if some of them might not be relevant. Since there is no direct information from the human about what is or is not relevant during a demonstration, and since the robot learns the task representation from even a single demonstration, it assumes that everything that it notices about the environment is important and represents it accordingly. In the next section we demonstrate how simple feedback cues can be used to signal to the robot the saliency of particular events. While this does not eliminate irrelevant environment state from being observed, it biases the robot to notice and (if capable) capture the key elements.

4.2. Learning in Environments With Distractors

The goal of the experiments presented in this section is to show the ability of the robots to learn from in environments with distractor objects, which are not relevant for the demonstrated tasks.

The task to be learned by the robot is similar to the moving objects task from above (Figure 10(a)): pick up the orange box placed near the light green target (the source), go through the “gate” formed by the yellow and light orange target, drop the box at the dark green target (the destination) and then come back to the source target. The orange and the yellow targets at the left are distractors that should not be considered as part of the task. In order to teach the robot that it has to pick up the box, the human led the robot to it and then, when sufficiently near it, placed it between the robot’s grippers. At the destination target, the teacher took the box from the robot’s grippers. Moments in time signaled by the teacher as being relevant to the task are: giving the robot the box while close to the *light green* target, teacher reaching the *yellow* and *light orange* target, taking the box from the robot while at the *green* target, and teacher reaching the *light green* target in the end. Thus, although the robot observed that it had passed the *orange* and distant *yellow* targets during the demonstration, it did not include them in its task representation, since the teacher did not signal any relevance while being at them.




Figure 10. The “Object manipulation” task in environments with distractors

We performed 10 human-robot demonstration experiments to validate the performance of our learning algorithm. In 9 of the 10 experiments the robot learned a structurally correct representation (sequencing of the relevant behaviors) and also performed it correctly. In one case, although the structure of the behavior network was correct, the learned values of one of the behavior’s parameters caused the robot to perform an incorrect task (instead of going between two of the targets the robot went to them and then around). The learned behavior network representation of this task is presented in Figure 11.




Figure 11. Task representation learned from human demonstration for the “Object manipulation” task

In Figure 10(b) we show the robot’s progress during the execution of the task, more specifically the instants of time or the intervals during which the postconditions of the behaviors in the network were true.

For the 9 out of 10 successes we have recorded, the 95% confidence interval for the binomial distribution of the learning rate is [0.5552 0.9975], obtained using a Paulson-Camp-Pratt approximation [2] of the confidence limits.

As a base-case scenario, to demonstrate the reliability of the learned representation, we performed 10 trials, in which a robot repeatedly executed one of the learned representations of the above task. In 9 of the 10 cases the robot correctly completed the execution of the task. The only failure was due to a time-out in tracking the green target.

5. Related work

The work presented here is most related to two important areas of robotics research: human-robot interaction and robot learning. Here we discuss its relation to both areas and state the advantages gained by combining the two in the context of adding social capabilities to agents in human-robot domains.

Most of the approaches to human-robot interaction so far rely on using predefined, common *vocabularies* of gestures [12], signs or words. These can be said to be using a *symbolic language*, whose elements explicitly communicate specific meanings. In this work, we show that communication between robots and humans can be achieved even without such explicit prior vocabulary sharing.

One of the most important forms of implicit communication, which has received a great deal of attention among researchers, is the use of various forms of body language. For example, it has been applied to humanoid robots (in particular head-eye systems), for communicating emotional states through face expressions [3] or body movements [4], where the interaction is performed through *body language*. While facial expressions are a natural means of interaction for a humanoid, or in general a “headed,” robot, they cannot be entirely applied to the domain of mobile robots, where the platforms typically have a very different, and non-anthropomorphic physical structure. In our approach, we demonstrate that the use of implicit, action-based methods for communicating and expressing intentions can be extended to the mobile robot domain, despite the structural differences between mobile robots and humans.

Teaching robots new tasks is a topic of great interest in robotics. In the context of behavior-based robot learning, methods for learning policies (situation-behavior mappings) have been successfully applied to single-robot learning of various tasks, most commonly navigation [7], hexapod walking [13] and box-pushing [14].

In the area of teaching robots by demonstration, also referred to as imitation, [10] demonstrated simplified maze learning, i.e., learning turning behaviors, by following another robot teacher. The robot used its own observations to relate the changes in the environment with its own forward, left, and right turn actions. [21] 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 in our lab is also exploring imitation based on mapping observed human demonstration onto a set of behavior primitives, implemented on a 20 DOF dynamic humanoid simulation [16, 11]. The key difference between the work presented here and those above is at the level of learning. The work above 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.

6. Conclusions

In this paper we presented an *action-based* approach to human-robot interaction and robot learning, both dealing with aspects of designing socially intelligent agents. The method was shown to be effective for interacting with humans using implicit, action-based communication and learning from experienced demonstration.

We argued that the means of communication and interaction of mobile robots which do not have anthropomorphic, animal, or pet-like appearance and expressiveness should not necessarily be limited to explicit types of interaction,

such as speech or gestures. We demonstrated that simple actions could be used in order to allow a robot to successfully interact with users and express its intentions. For a large class of intentions such as: *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 thus obtain assistance.

We also presented a methodology for learning from demonstration in which the robot learns by relating the observations to the known effects of its behavior repertoire. This is made possible by our behavior architecture that has a perceptual component (*abstract behavior*) which embeds representations of the robot's behavior goals. We demonstrated that the method is robust and can be applied to a variety of tasks involving the execution of long, and sometimes even repeated sequences of behaviors.

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

7 Acknowledgments

This work is 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.

References

- [1] R. C. Arkin. *Behavior-Based Robotics*. MIT Press, CA, 1998.
- [2] C. R. Blyth. Approximate binomial confidence limits. *Journal of the American Statistical Association*, 81(395):843–855, September 1986.
- [3] C. Breazeal and B. Scassellati. How to build robots that make friends and influence people. In *Proc., IROS, Kyonju, Korea*, pages 858–863, 1999.
- [4] L. D. Canamero and J. Fredslund. How does it feel? emotional interaction with a humanoid lego robot. Tech Report FS-00-04, AAAI Fall Symposium, 2000.
- [5] A. David and M. P. Ball. The video game: a model for teacher-student collaboration. *Momentum*, 17(1):24–26, 1986.
- [6] D. C. Dennett. *The Intentional Stance*. MIT Press, Cambridge, 1987.
- [7] M. Dorigo and M. Colombetti. *Robot Shaping: An Experiment in Behavior Engineering*. MIT Press, Cambridge, 1997.
- [8] S. T. et al. A second generation mobile tour-guide robot. In *Proc. of IEEE, ICRA*, 1999.
- [9] T. M. et al. An office conversation mobile robot that learns by navigation and conversation. In *Proc., Real World Computing Symp.*, pages 59–62, 1997.
- [10] G. Hayes and J. Demiris. A robot controller using learning by imitation. In *Proc. of the Intl. Symp. on Intelligent Robotic Systems*, pages 198–204, Grenoble, France, 1994.
- [11] O. C. Jenkins, M. J. Matarić, and S. Weber. Primitive-based movement classification for humanoid imitation. In *Proc., First IEEE-RAS Intl. Conf. on Humanoid Robotics*, Cambridge, MA, MIT, 2000.
- [12] D. Kortenkamp, E. Huber, and R. P. Bonasso. Recognizing and interpreting gestures on a mobile robot. *Proc., AAAI*, pages 915–921, 1996.
- [13] P. Maes and R. A. Brooks. Learning to coordinate behaviors. In *Proc., AAAI*, pages 796–802, Boston, MA, 1990.
- [14] S. Mahadevan and J. Connell. Scaling reinforcement learning to robotics by exploiting the subsumption architecture. In *Eighth Intl. Workshop on Machine Learning*, pages 328–337, 1991.
- [15] M. J. Matarić. Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2–3):323–336, 1997.
- [16] M. J. Matarić. Sensory-motor primitives as a basis for imitation: Linking perception to action and biology to robotics. In C. Nehaniv and K. Dautenhahn, editors, *Imitation in Animals and Artifacts*. MIT Press, Cambridge, to appear, 2001.
- [17] F. Michaud and C. S. Roball - an autonomous toy-rolling robot. In *Proc. of the Workshop on Interactive Robotics and Entertainment*, 2000.
- [18] M. N. Nicolescu and M. J. Matarić. Experience-based representation construction: learning from human and robot teachers. In *Proc., IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pages 740–745, Maui, Hawaii, USA, Oct 2001.
- [19] M. N. Nicolescu and M. J. Matarić. A hierarchical architecture for behavior-based robots. In *Proc., First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Bologna, ITALY, July 2002.
- [20] B. Scasellati. Investigating models of social development using a humanoid robot. In B. Webb and T. Consi, editors, *Biorobotics*. MIT Press, to appear, 2000.
- [21] S. Schaal. Learning from demonstration. In M. Mozer, M. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems 9*, pages 1040–1046. MIT Press, Cambridge, 1997.
- [22] B. B. Werger. Ayllu: Distributed port-arbitrated behavior-based control. In *Proc., The 5th Intl. Symp. on Distributed Autonomous Robotic Systems*, pages 25–34, Knoxville, TN, 2000. Springer.