

## Autonomy and Supervision for Robot Skills and Tasks Learned from Demonstration

### Odest Chadwicke Jenkins

Interaction Lab, Robotics Research Lab  
Computer Science Department  
University of Southern California  
941 W. 37th Place, SAL 227  
Los Angeles, CA, USA 90089-0781  
cjenkins@usc.edu

### Monica N. Nicolescu

Department of Computer Science  
and Engineering, MS 171  
University of Nevada-Reno  
1664 N. Virginia St.  
Reno, NV, USA 89557  
monica@cs.unr.edu

### Maja J Matarić

Interaction Lab, Robotics Research Lab  
Computer Science Department  
University of Southern California  
941 W. 37th Place, SAL 227  
Los Angeles, CA, USA 90089-0781  
mataric@usc.edu

### Introduction

One of the main goals of robotics is to assist or replace people performing tedious or dangerous tasks. Typically, robot controllers for such tasks are designed and implemented by specialists, and the process requires hand-crafting a new and different controller for each particular task. Learning from demonstration (Schaal 1997) and imitation (Matarić 2002; Schaal 1999) offer alternatives to hand-crafting of robot controllers, emphasizing collaborative human-robot interaction as a natural and accessible method for robot programming.

Demonstration-based learning and supervisory control have certain complementary properties. Both approaches maintain human presence in the robot's control process. Supervisory methods do so by yielding robot autonomy to human control. In contrast, demonstration-based methods allow for robot autonomy; in those, the human does not directly operate the robot, but provides teaching and feedback. Another complementary property of the two methods involves the amount of communication between the human and the robot. Potential human-robot interaction modalities, such as speech and gesture, may not be capable of conveying information at a sufficient rate to maintain direct human tele-operation of the robot. Furthermore, complex control problems are not always amenable to or practical for tele-operation, due to the dimensionality of the control space and the required skill and concentration on the part of the human operator.

The above scenarios can all benefit from a combination approach, with the robot maintaining autonomy over low-level control while receiving direct high-level control from a human operator or high-level feedback from a human teacher. Conversely, having the human in the loop at various levels of control is also of value in demonstration-based learning, as it is a subjective process susceptible to errors in the demonstration, observation, and generalization.

In this paper, we discuss the potential of such combination approaches to learning low-level robot skills and high-level robot task from demonstration, intended as a foundation for higher-level supervised control and collaboration in a variety of interaction modalities. In particular, we discuss a com-

bined *exemplar-based representation* of skills (e.g., Jenkins & Matarić (2003) and Jenkins & Matarić (2004)) and tasks (e.g., Nicolescu & Matarić (2002)), a method for their unsupervised extraction from demonstration, and supervised refinement through human feedback and operation. We place the discussion in the context of three scenarios for supervisory control, along with potential applications and avenues for future work.

### Representing Skills and Tasks Learned From Demonstration

For truly autonomous robots, an effective method for learning from demonstration should be able to uncover descriptions of skills and tasks contained in the demonstration data. We have recently proposed methods for unsupervised extraction of skills from human motion (Jenkins & Matarić 2003) and tasks from human demonstration to a mobile robot (Nicolescu & Matarić 2001). These methods were validated on vastly different robot platforms, but they share a common representation amenable to autonomous and supervised control.

In our representation, a *skill behavior* is defined by a set of exemplars in the configuration space of the robot and a *task behavior (or abstract behavior)* is a task description that coordinates those skills in a goal-directed fashion. As described in Jenkins & Matarić (2003), unlabeled and unsegmented motion data from demonstrations can be automatically partitioned into segments, clustered into distinct behaviors, and generalized into a repertoire of skill behaviors. The resulting behaviors are represented in a form that can be used not only for control but also for prediction. This predictive capability enables the skill behaviors to provide expected future robot configurations useful for specifying control desires or for matching against observed demonstrator configurations.

Task or abstract behaviors provide purposeful, goal-directed robot control based on the foundation of skill behaviors described above. In Nicolescu & Matarić (2003), an approach to learning task behaviors from demonstration is described that assumes a predefined set of skill behaviors, complementary to Jenkins & Matarić (2003). In that approach, skill behaviors are used to partition and classify (through matching observations to skill postconditions) an

individual demonstration into a sequential task description of the form  $A \rightarrow B \rightarrow C$  or  $A \rightarrow D \rightarrow C$ . Task descriptions across multiple demonstrations are then generalized to produce a single description that retains the common characteristics through specific-to-general learning. For instance, the general description  $A \rightarrow (B \text{ or } D) \rightarrow C$  is produced from the two examples above. Because the task description is encoded using a set of skill behaviors as controllers, the resulting generalized task description can be used to direct the skill controllers to perform the task on the robot independent of a particular configuration or situation. Additionally, the task behaviors can be structured hierarchically to construct new, higher-level composite task behaviors.

## Relating Skills and Tasks

While sharing the exemplar-based philosophy and the use of such *behavior primitives* (Mataric 2002), skill representations utilized by Nicolescu & Mataric (2002) and Jenkins & Mataric (2003) differ in their form, at least in part due to the type of robot platform on which they were employed: mobile robots and humanoids, respectively. In particular, the following issues require further attention:

- *Perceptual-motor versus Sensory-motor primitives as skills.* The perceptual-motor skills used by Jenkins and Mataric assume that data gathered through observation have been preprocessed and transformed into the proprioceptive (internal) configuration space of the robot. In contrast, sensory-motor skills used by Nicolescu and Mataric use basic sensory preprocessing and consider exteroceptive factors, the robot's external state for the environment.
- *Preconditions and Postconditions.* Given their consideration of exteroceptive factors, Nicolescu and Mataric require explicit preconditions and postconditions for each skill. In contrast, preconditions and postconditions are implicit in the skills of Jenkins and Mataric and would require explicit extraction.

Our continuing work is exploring these issues.

## Scenarios for Supervisory Control

Next, we discuss situations where supervisory control/feedback is necessary or beneficial for the representation of the described skill and task behaviors. We also describe avenues for future research in the context of these scenarios.

### Supervisory Feedback for Autonomous Task Control and Refinement

As described in Nicolescu & Mataric (2003), task behaviors whose descriptions were extracted/learned from demonstration, can be used for autonomously controlling a robot to perform the demonstrated task. This autonomy is enabled by having a repertoire of underlying skills that provide the robot with autonomous control and the ability to classify the observed demonstration into those skills and parametrize and execute them accordingly.

Given multiple demonstrations of a specific task, the learned generalized task behavior encodes the specifics of

each demonstration while incorporating their common components. Consequently, the generalized task behavior is capable of autonomous robot control of the given task in a variety of related situations.

In a significant number of cases, however, a task behavior may not perform as intended by the human during autonomous execution due to factors such as sparsity in the input set and errors in generalization. In such cases, the human intervenes to provide feedback to the robot, with the result being the refinement of the robot's task description.

As described, this method allows for refinement of task descriptions but not their underlying repertoire of skills. Ideally, both the task and its underlying skills should be refinable based on user feedback. Skills used in Nicolescu & Mataric (2003) are not readily modifiable due to their manually coded nature, but the combined exemplar-based skill and task representation presented in the previous section could remedy this problem. However, challenges remain, including:

- **Determining what should be refined.** Manual and automated methods for determining whether refinement should take place for the task descriptions or for an individual skill are needed. In their absence, wholesale combined refinement of a task and underlying skills requires increased effort from the human supervisor. Additionally, such wholesale refinement requires a means of adaptive capable of retaining the desirable while discarding the undesirable properties of the skills.
- **Avoiding task specificity in skills.** A known danger in refining skills for a particular task is overfitting, i.e., convergence towards the single task. Avoiding such overfitting is a recognized problem in learning.

### Collaborative High-level Supervisory Control

As discussed in the previous subsection, Nicolescu & Mataric (2003) utilizes a formal relationship between the robot and the human input through speech and gesture communication. This method is suited for robots learning from humans and other robots and for autonomous performance by a robot. For certain tasks, however, a more collaborative relationship is needed between the robot and the human, or among multiple robots. Robot soccer is a popular example of multi-robot collaboration, where robots perform autonomously but not independently; instead, they are indirectly supervised by others on the team via sensing and communication. In addition, they may (at some future time) play with human teammates who use speech and gesture for interaction. Such collaborations assume a significant amount of *a priori* domain knowledge. Methods for learning from demonstration could provide an excellent means for acquiring such collaborative tasks and skills without the intractability of exploring vast learning spaces (as is the case in robot teams) or the specificity of hardcoded controllers and interactions.

Collaborating at this level involves major challenges, including:

- **How to observe collaborative demonstration?** It is difficult to observe complex interactions and collaboration

without detailed *a priori* models of expected interactions.

- **How to segment, cluster, and generalize collaborative demonstration?** In particular, it is an open question as to how to generalize a robot's interactions with collaborators and adversaries.
- **How to model other collaborators?** We assume collaborators will not be able to provide detailed communications. Thus, collaborator interactions will require some type of a generalizable model.

### Direct High-Level Supervisory Control

In contrast to the previous two scenarios, autonomous task-level control may not be necessary when a human can continually provide high-level tele-operation. In such situations, the human directs a repertoire of low-level skills for controlling the robot's actuation system. One application of this approach is in *brain-machine interfaces for neural prostheses*. As described by Black *et al.* (2003), devices for direct interfaces are advancing toward providing greater ability to help people with motor impairments. In the foreseeable future, however, such devices will become bandwidth-limited; the limited communication channels could be complemented by a repertoire of exemplar-based skills, such as those described by Jenkins & Matarić (2003). Human users could use high-level sparse control to direct autonomous motor programs encoded as skills that actuate a robot or some other relevant platform.

Such direct high-level control brings up several questions, including:

- **What are appropriate vocabularies of skills?** This issue is addressed in part by the automated skill derivation method provided by Jenkins & Matarić (2003). However, that method is dependent upon having demonstration data as input, which may not be unavailable.
- **How can signals generated from human high-level operation be reliably mapped onto a repertoire of skills?** Classification of observed human instruction onto the skill repertoire for non-trivial vocabularies and real activity perception remains a challenging problem.
- **How can we extract structure from human high-level operation to build increasingly rich hierarchical sets of autonomous behaviors?** Such task behaviors would expand the ability of the human operator toward new and more complex tasks.

In summary, we have discussed the potential benefits and challenges of learning low-level robot skills and high-level robot task from demonstration/by imitation. Such skills and tasks can serve as a foundation for higher-level supervised control and collaboration in a variety of interaction modalities. We have used our past work as an example of promising avenues and outstanding research problems.

### References

- Black, M. J.; Bienenstock, E.; Donoghue, J. P.; Serruya, M.; Wu, W.; and Gao, Y. 2003. Connecting brains with machines: The neural control of 2d cursor movement. In *IEEE/EMBS Conference on Neural Engineering*, 580–583.
- Jenkins, O. C., and Matarić, M. J. 2003. Automated derivation of behavior vocabularies for autonomous humanoid motion. In *Autonomous Agents and Multiagent Systems (AAMAS 2003)*, 225–232.
- Jenkins, O. C., and Matarić, M. J. 2004. Performance-derived behavior vocabularies: Data-driven acquisition of skills from motion. *Submitted to the International Journal of Humanoid Robotics*.
- Matarić, M. J. 2002. Sensory-motor primitives as a basis for imitation: Linking perception to action and biology to robotics. In Nehaniv, C., and Dautenhahn, K., eds., *Imitation in Animals and Artifacts*. MIT Press. 392–422.
- Nicolescu, M. N., and Matarić, M. J. 2001. 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* 31(5):419–430.
- Nicolescu, M., and Matarić, M. J. 2002. A hierarchical architecture for behavior-based robots. In *Autonomous Agents and Multi-Agent Systems (AAMAS 2002)*, 227–233.
- Nicolescu, M. N., and Matarić, M. J. 2003. Natural methods for robot task learning: Instructive demonstration, generalization and practice. In *Proc., Second Intl. Joint Conf. on Autonomous Agents and Multi-Agent Systems*, 241–248.
- Schaal, S. 1997. Learning from demonstration. In Mozer, M.; Jordan, M.; and Petsche, T., eds., *Advances in Neural Information Processing Systems 9*. The MIT Press. 1040–1046.
- Schaal, S. 1999. Is imitation learning the route to humanoid robots. *Trends in Cognitive Sciences* 3(6):233–242.