MOTOR SCHEMAS IN ROBOT LEARNING

Lynne E. Parker

Department of Electrical Engineering and Computer Science, University of Tennessee

Knoxville, Tennessee

USA

leparker@utk.edu

Synonyms

Movement primitives in robot learning; Macro actions in robot learning; Basis behaviors for robot learning.

Definition

Motor schemas used for *robot learning* are sequences of action that accomplish a goal-directed behavior, or a task. Motor schemas in robot learning are also known as *movement primitives*, *basis behaviors*, *units of action*, and *macro actions*. Rather than representing the simplest elementary actions available to the robot, such as a simple command to a robot actuator, schemas and motion primitives represent a higher-level abstraction of robot actions, such as "avoid obstacles", "wander", "walk", "grasp a cup", and "move to goal". These schemas and motion primitives define control policies that are encoded with only a few parameters, and serve as the *basis set*, or *movement vocabulary*, of the robot. Such primitives are sufficient for generating the robot's entire repertoire of motions via the combination of schemas or primitives. The schema can serve as a basis for robot learning, since it provides an abstraction that can be represented with fewer parameters, thus reducing the complexity of robot learning. This reduction in learning complexity allows robot learning to scale to more complex robots or tasks, thus making practical applications tractable.

Theoretical Background

The use of motor schemas in robotics first became popular in the 1980's, especially with the work of Lyons and Arbib (1989), and of Arkin (1987). Their development of the motor schema concepts was inspired by similar concepts in psychology and the neurobiological sciences. These early researchers recognized that ideas for how motor behavior control is achieved in animals (e.g., frogs, as studied by Arbib) or in humans can provide a model for how similar behaviors can be created in robots. As defined in this early work, a robot schema consists of a list of input and output ports, a local variable list, and a behavior that defines how the input is processed to generate the output. Robot schemas can be of two broad types – *perceptual schemas* and *motor schemas*. Perceptual schemas, which can be embedded inside motor schemas, process input from environmental sensors on the robot to provide information to motor schemas. The motor schemas then generate output control vectors that represent the way the robot should move to achieve a goal, in response to the perceived stimuli. Schemas are independent, and can run concurrently with other schemas. A network of schemas can be built by manually connecting the outputs of one schema to the inputs of another. The output from multiple motor schemas can be combined using techniques from potential fields, such as vector addition. Motor schemas can be grouped to form more complex behaviors, which are sometimes called *behavior assemblages*. At the higher level, a nested network is established to represent collaboration among multiple robots.

In more recent years, roboticists have made less use of the *schema* terminology, preferring instead to describe robot design in terms of *movement primitives*, *basis behaviors*, *units of action*, or *macro actions*. These latter terms still capture many of the same ideas as the robot schema concept, although the implementation and realization of the movement primitives may be somewhat different from the original schema concept.

While human designers of robot systems can try to manually define the specific motor schemas and their inter-connections that will be used in a robot system to solve a given task, this manual design process proves to be quite difficult for most practical applications. The difficulties arise in (1) the inability to anticipate the interactions of multiple schemas (or movement primitives), (2) the inability to discover the proper schema (primitive) combinations that achieve the required task, and (3) the unexpected interactions of the robot with the environment in which it operates. Because of these difficulties, learning approaches are preferred that enable the robot to learn and adapt its behavior from the fundamental behavior building blocks (i.e., schemas and/or movement primitives) provided by the human designer. Schemas and/or movement primitives are especially helpful in this context, as they provide the robot with fundamental building blocks that can be combined and parameterized as appropriate to achieve the task at hand. Learning techniques allow the robot to learn the appropriate sequences of schemas/behaviors that will accomplish the task, or the appropriate parameters with which to instantiate the schemas/behaviors, without requiring the designer to fully specify how the task should be accomplished. The overall rate of learning for a task has been shown to be increased by breaking down the task into subtasks, then learning at the subtask level, rather than monolithically at the higher-task level. Thus, the schema-based approach to robot learning provides a helpful abstraction for making the learning task achievable.

Important Scientific Research and Open Questions

A primary challenge in the use of motor schemas/primitives in robot learning is determining how to select, parameterize, sequence, or combine the predefined schemas or primitives to achieve a given task. One approach to this challenge is to have a human teacher or trainer to illustrate the desired task; the robot then seeks to emulate this demonstrated task. Much current research investigates this idea of robot learning via imitation of human actions, by building up from existing motion primitives (e.g., Breazeal and Scassellati, 2002). In this approach, the robot must observe the human and determine which of the human's actions are relevant for the current task. This challenge includes the 3D perception problem of perceiving human movement through vision, as well as the attention problem, in which the robot selectively focuses on the aspects of the motion that are particularly relevant to the task to be learned. Once the action has been perceived, the robot must transform the perception into its own motor actions that achieve the same result.

However, it is not trivial for a robot to determine which motor schemas, or motion primitives, correspond to the demonstrated task. One approach (Schaal, 2003) is to execute each motor primitive, observe its outcome, and evaluate the result using a performance criterion that compares the similarity between the teacher's behavior and the robot's generated behavior. Another approach makes use of predictive forward models, in which each movement primitive tries to predict the next observed motion, based on the current state of the teacher. The motion primitive with the best prediction capabilities would be selected as the best match.

Another way of mapping the behavior of the trainer onto the robot's existing repertoire of basic/primitive capabilities has been proposed by Nicolescu, et al. (2008). This work defines a behavior-based approach to learning from demonstration that uses behavior fusion to provide bottom-up generalization to new situations. This approach learns a coordination policy that linearly fuses the combined output of preexisting robot behaviors, which are expressed as schemas or potential fields, in a manner that matches the teacher's demonstration. The learning of this coordination is expressed as a fusion estimation problem, i.e., state estimation in the space of linear combinations of primitive behaviors. For domains such as mobile robotics, fusion estimation is often subject to ambiguous changes in world state that are attributable to a large space of solutions. To account for this ambiguity and dynamic changes to the user's fusion policy, a particle filter is used to infer fusion estimates from robot sensory observations and motor commands. This learning technique allows for learning of superposition behavior fusion from existing innate robot primitives, and learning of sequential activities from multiple superposition fusion primitives.

Another approach for addressing this challenge is the work by Maja Mataric, which is based on the discovery of "mirror neurons" in monkeys, which fire when the monkey both observes a goal-oriented action, and when it performs the same action. The entire approach to robot imitation learning combines several cognitive approaches, including movement perception through a specialized selective attention system, direct sensory-motor mapping between the perceived and executable movement, movement generation through a system of composable motor primitives, and learning of new movements and skills by building on existing repertoire of motor primitives through classification and combination. In this work, three types of motor primitives are defined: discrete straight-line movements, oscillatory movements, and postural movements that define large subsets or whole-body arrangements of joints. Learning techniques such as reinforcement learning can be used to parameterize these primitives appropriately. In related work, certain types of motor primitives can also be learned by tracking the movements of humans, using Principal Components Analysis to extract the most relevant features from the motion data, and then using these features to reconstruct the original movement on the robot.

An alternative approach for robot learning of more complex tasks from primitive schemas is to enable the robot itself to explore its capabilities, rather than following the guidance of a human trainer. This type of approach is often called *constructivist* robot learning, which is a method for learning new knowledge and skills based upon past experience. This type of learning is recognized to be a common method used by humans from infancy to adulthood for lifelong learning. Because much of human learning seems to be based on schema building blocks, a similar approach is used in robotic applications. For example, Gary Drescher, as well as Harold Chaput, both developed schema-based constructivist learning models to computationally emulate an infant exploring the environment using very basic perceptual schemas and motor schemas. Their work concentrated on the biological verification of the constructivist point of view using very basic level schemas that reflect the inherent abilities of an infant.

A related approach to schema learning that does not involve a human teacher is the work of Tang and Parker (2008), who developed the SB-CoRLA (for Schema-Based, Constructivist Robot Learning Architecture) architecture, in which robots are able to build up combinations of schemas, called "chunks", which can then be used to improve the robot's efficiency in performing future tasks. The approach involves both an offline learning phase and an online learning phase. In the offline learning phase, which occurs when the robot is not busy performing tasks, the robot uses an evolutionary search technique to analyze its schema repository for highly fit partial solutions to tasks of interest to the robot. These solutions are then saved as chunks for future use in the online phase. In the online learning phase, the robot uses both the individual schemas and the schema chunks to quickly find good solutions for addressing the task at hand.

A unifying theme of all these approaches is the recognition that the use of fundamental building blocks, in the form of schemas, motion primitives, basis behaviors, etc., is a powerful way to make the robot learning problem tractable. By properly defining motor schemas for a given application, developed techniques can be used to select, parameterize, sequence, or combine the predefined schemas or primitives to enable the robot to achieve a given task. Many open issues remain, however. Certainly, more research is needed to deal with the perceptual understanding of the effects of motions, whether motions generated by human teachers or by the robot itself. Further, it is still currently difficult for robots to understand the high-level goals or objectives of demonstrated movement, and to determine how to best map these to the predefined repertoire of motion primitives. Additional open challenges include determining the appropriate set of schemas for a given application, and determining how to enable a robot to learn new schemas, in order to build up the available repertoire of motor schemas.

Cross-References

→ Action schemas

→ Learning action affordances and action schemas

→ Motor schemas

→ Schemas

→ Schema-based architectures of machine learning

→ Schema-based learning

→ Sensori-motor schemas

→ Robot learning

→ Developmental robotics

References

Arkin, R.C. (1987). Motor schema based navigation for a mobile robot: An approach to programming by behavior. *Proceedings of the IEEE International Conference on Robotics and Automation*.

Breazeal, C. & Scassellati, B. (2002). Robots that imitate humans. *TRENDS in Cognitive Sciences*, 6 (11), 481-487.

Lyons, D. M. & Arbib, M. A. (1989). A formal model of computation for sensory-based robotics, *IEEE Transactions on Robotics and Automation*, 5 (3), 280-293.

Nicolescu, M, Jenkins, O. C., Olenderski, A., & Fritzinger, E.. (2008). Learning behavior fusion from observation, *Interactive Studies Journal*, special issue on Robot and Human Interactive Communication, 9 (2), 319-352.

Schaal, S., Ijspeert, A., & Billard, A. (2003). Computational approaches to motor learning by imitation, *Philosophical Transactions of the Royal Society of London: Biological Sciences*, 358 (1431), 537-547.

Tang, Y. & L. E. Parker (2008). Towards schema-based, constructivist robot learning: Validating an evolutionary search algorithm for schema chunking. *Proceedings of the IEEE International Conference on Robotics and Automation*.