Machine Learning for Motor Skills in Robotics

Jan Peters

Autonomous robots that can adapt to novel situations has been a long standing vision of robotics, artificial intelligence, and the cognitive sciences. Early approaches to this goal during the heydays of artificial intelligence research in the late 1980s, however, made it clear that an approach purely based on reasoning or human insights would not be able to model all the perceptuomotor tasks of future robots. Instead, new hope was put in the growing wake of machine learning that promised fully adaptive control algorithms which learn both by observation and trial-and-error. However, to date, learning techniques have yet to fulfill this promise as only few methods manage to scale into the high-dimensional domains of manipulator and humanoid robotics and usually scaling was only achieved in precisely pre-structured domains. We have investigated the ingredients for a general approach to motor skill learning in order to get one step closer towards human-like performance. For doing so, we study two major components for such an approach, i.e., firstly, a theoretically well-founded general approach to representing the required control structures for task representation and execution and, secondly, appropriate learning algorithms which can be applied in this setting.

1 Introduction

Despite an increasing number of motor skills exhibited by manipulator and humanoid robots, the general approach to the generation of such motor behaviors has changed little over the last decades [4]. The roboticist models the task as accurately as possible and uses human understanding of the required motor skills in order to create the desired robot behavior as well as to eliminate all uncertainties of the environment. In most cases, such a process boils down to recording a desired trajectory in a pre-structured environment with precisely placed objects. If inaccuracies remain, the engineer creates exceptions using human understanding of the task. While such highly engineered approaches are feasible in well-structured industrial or research environments, it is obvious that if robots should ever leave factory floors and research environments, we will need to reduce the strong reliance on hand-crafted models of the environment and the robots exhibited to date. Instead, we need a general approach which allows us to use compliant robots designed for interaction with less structured and uncertain environments in order to reach domains outside industry. Such an approach cannot solely rely on human knowledge but instead has to be acquired and adapted from data generated both by human demonstrations of the skill as well as trial and error of the robot.

The tremendous progress in machine learning over the last decades offers us the promise of less human-driven approaches to motor skill acquisition. However, despite offering the most general way of thinking about data-driven acquisition of motor skills, generic machine learning techniques, which do not rely on an understanding of motor systems, often do not scale into the domain of manipulator or humanoid robotics due to the high domain dimensionality. Therefore, instead of attempting an unstructured, monolithic machine learning approach to motor skill aquisition, we need to develop approaches suitable for this particular domain with the inherent problems of task representation, learning and execution addressed separately in a coherent framework employing a combination of imitation, reinforcement and model learning in order to cope with the complexities involved in motor skill learning. The advantage of such a concerted approach is that it allows the separation of the main problems of motor skill acquisition, refinement and control. Instead of either having an unstructured, monolithic machine learning approach or creating hand-crafted approaches with pre-specified trajectories, we are capable of acquiring skills, represented as policies, from demonstrations and refine them by trial and error. Using learning-based approaches, we can achieve accurate control without accurate analytical models of the complete system.

2 Motor Skill Learning

Our principal objective is establishing the foundations for a general framework for representing, learning and executing motor skills for robotics. As can be observed from this question, this goal requires three building blocks, i.e., appropriate representations for movements, learning algorithms which can be applied to these representations and a learned transformation that allows the execution of the kinematic policies in the respective task space on robots.

2.1 Important Aspects of Motor Skills

We briefly address the three essential aspects of motor skill learning, i.e., representation, learning and execution, in this section.

Representation. For the representation of motor skills, we can rely on the insight that humans, while being capable of performing a large variety of complicated movements, restrict themselves to a smaller amount of primitive motions. As suggested by Ijspeert et al. [2], such primitive movements can be represented by nonlinear dynamic systems. We can