Learning Algorithms for Physical Systems: Challenges and Solutions

Ion Matei, Palo Alto Research Center (PARC)

Giving a machine the ability to learn, adapt, or repair itself are among the initial and most ambitious goals of computer science. A learning process uses information obtained during one interaction with its environment to improve its performance during future interactions. We focus on two learning results: models and policies. Models help us to make sense of our world. They provide a framework or structure to help us understand a large or complex concept. They enable us to analyze it, discover its weakness and improve it. Policies are a set of actions that enable a system achieving a set of objectives. Learning processes are enabled by algorithms such as machine learning algorithms. Machine learning algorithms are effective tools in a variety of applications in signal and information processing applications. Such applications usually benefit from large amounts of data. Systems analytics tasks such as control, diagnosis or prognostic can also benefit from machine learning. Unfortunately it is often the case that large amounts of data are not available for this type of tasks. In this talk I will discuss two problems: how to learn policies for physical systems with little experimental data, and how to build acausal models for physical systems.

Learning models of physical systems or learning policies for optimizing some objective function relies on potentially large data-sets that describe the behavior of the system. When such data-sets are not available, we look for alternatives to supplement the training data set. Such alternatives may include synthetic data originating from simulations, when a model of the system is available, or from experiments on the real system. Both alternatives bring their own challenges. For example, to build a physics-based model from first principles, we need information about the physical processes that govern the behavior of the system, and the set of parameters that control these physical processes. Unfortunately, it is usually the case that such parameters are not easily obtained. Often the components of a physical system originate from different manufacturers who are typically not keen to share technical proprietary information about their products. Performing experiments in the field may also not be feasible, since the operator of the system may not be willing to disturb the normal operation of deployed systems. We propose a third alternative that neither involves model simulations, nor performing experiments.

Our approach is based on using available data as a seed for generating a new source of synthetic data that describes the behavior of the system. In particular, we apply a transformation to a time series describing a trajectory of the system, and obtain a different time series that describes another feasible trajectory. This is called symmetry, and it represents a function that maps a trajectory to one or many feasible trajectories. The number of additional trajectories depends on the symmetry type. In the case of discrete symmetries, one additional trajectory can be generated. In the case of parameterized symmetries, theoretically, we can generate infinitely many. Finding symmetries is a rather challenging task. In the case of physical systems, we can look for possible symmetries such as rotation, translation or scaling. Checking whether a map is a symmetry requires some knowledge about the system. We can check if a map is a symmetry by using qualitative models of physical systems induced by the physical laws governing their behavior. By leveraging the geometric properties of physical systems such as symmetries we demonstrate how to improve reinforcement algorithms used for policy learning to deal with limited experimental data.

Machine learning algorithms usually produce input-output maps that hide information about the physical system that generated the data. It other words we lose explainability. This property is very significant in diagnosis and prognostics applications. Explainability is preserved though in physics based model. In the second part of the talk we show how we can learn acausal models of physical components in partially known systems. An acausal system is composed of variables attached to its components and relations between them. The relations are induced by the parameterized constitutive equations and the connections between components. The parameters of the constitutive equations are usually constrained within some feasibility set. Our approach is based on the following steps: (i) we select a parametric mathematical model for the constitutive equations; (ii) we learn the parameters and a representation/model of the parameter constraints. For the latter, we use an explore-exploit strategy, where in the beginning we focus on exploration to learn the constraints representation. Later, as the constraints representation becomes more informative, the focus is shifted towards finding the best feasible parameters of the constitutive equations. We use machine learning inspired models such as classifiers to represent the feasibility constraints. Feasible component models ensure simulatability of the overall system model. This is a necessary condition for learning the component parameters. It also ensures reusability: when the model of the physical component is used in a different configuration, the real and simulated behaviors must be similar. An alternative to constraint learning is proposing parametrized components model that are inherently feasible. We use the passivity notion to generate templates for feasible component models and constraints for their parameters. Having a representation for parameter constraints serves to: decrease the complexity of the search; provide a good initial condition when new data is available for refining the component's parametric model and the constraints representation; enable a physical interpretation of the component model.

Keywords:

Physical systems: entities that can be separated from their environment by means of conceptual limits. They interact with their environment, which results in observable changes over the time.

Models: representations of systems used to answer questions via analysis and simulation. They can describe physical, biological or information systems.

Acausal models: models (abstractions) of physical systems. They are composed of variables and relations between them. The variables are functions of time to observable quantities. The relations act as constraints between the values variables take at each instant Models interact through constraints between some of their variables.

Symmetries: maps that given a trajectory of the system variables produce another trajectory.

Reinforcement learning: area of machine learning concerned with how software agents ought to take actions in an environment so as to maximize some notion of cumulative reward.