

Statistical learning & dynamical systems: exploiting hidden low-dimensional structures

Mauro Maggioni *

mauro.maggioni@jhu.edu

Inferring the laws of motion of physical systems from observations is a fundamental challenge. Many different tools have been brought to bear in different scenarios, with statistical and machine learning techniques becoming more prominent and useful with the abundance of data. Many challenges still remain. In this talk I discuss two examples of geometry-based statistical learning techniques for learning approximations to certain classes of high-dimensional dynamical systems. In the first scenario, we consider systems that are well-approximated by a stochastic process of diffusion type on a low-dimensional manifold. Neither the process nor the manifold are known, but we assume we have access to a way of sampling initial conditions and to a (typically expensive) simulator that can return short paths of the stochastic system, given an initial condition. We introduce a statistical learning framework for estimating local approximations to the system, for stitching these pieces together and form a fast global reduced model for the system, called ATLAS. ATLAS is guaranteed to be accurate in the sense of producing stochastic paths whose distribution is close to that of paths generated by the original system not only at small time scales, but also at very large time scales (under suitable assumptions on the dynamics). We discuss applications to homogenization of rough diffusions in low and high dimensions, as well as systems with separation of time scales. In the second scenario we consider a system of interacting agents: given only observed trajectories of the system, we are interested in estimating the interaction laws between the agents. We consider both the mean-field limit (i.e. the number of agents going to infinity) and the case of a finite number of agents, with an increasing number of observations. We show that at least in particular cases, where the interaction is governed by an (unknown) function of distances, the high-dimensionality of the state space of the system does not affect the learning rates. We prove that in these case in fact we can achieve an optimal learning rate for the interaction kernel, equal to that of a one-dimensional regression problem. We exhibit efficient algorithms for constructing our

*Department of Mathematics, Johns Hopkins University, 3400 N Charles St, 405 Krieger Hall, Baltimore, MD 21218, USA

estimator for the interaction kernels, with statistical guarantees, and demonstrate them on various examples.