Discovery and Learning in Predictive State Representations

Adrian Ghizaru

Predictive state representations are an emerging tool for modelling dynamic systems, and were introduced in 2001 by Satinder Singh, Michael Littman and Richard Sutton. Unlike partially observable Markov decision processes, which are today's standard solution to this task, PSRs do not postulate the existance of hidden states (that are often semantically ambiguous). Instead, they derive the probabilities of system trajectories by combining the probabilities of certain sequences of action-observation pairs. This is quite similar to the way that the concepts of 'basis' and 'span' are defined in linear algebra.

My project this summer was aimed at laying the foundation for various empirical studies of the properties of PSRs. As such, I began implementing an algorithm outlined by the proponents of PSRs. It is meant to discover what the base sequences are solely through interaction with the dynamic system. Once these "core tests" are found, a second phase of the algorithm estimates the parameters according to which these tests are to be combined to obtain the system model. However, at the time of this writing this second phase is yet to be implemented.

During my presentation, I will explain more exactly the theory of PSRs and I will detail the workings of my implementation of the discovery algorithm mentioned above.