YUNUS EMRE DEMIRCI, Queen's University

On Regularity and Ergodicity of Partially Observable Markov (Decision) Processes

In this talk, we study time-homogeneous hidden Markov models where states are not directly observable, also known as partially observable Markov processes. Instead, these states are observed through a measurement channel. The initial state is determined by an initial distribution, and as new observations are made, we update the state's conditional probability measure given the measurements, leading to a nonlinear filtering process.

The focus of our study is on the regularity properties for these nonlinear filters under the Wasserstein metric. We present conditions which lead to geometric ergodicity, implying that the filter process converges to invariance at an exponential rate (as a probability measure valued Markov chain). While unique ergodicity of such filter processes had been studied in the literature, such a geometric ergodicity result appears to be new. We also provide complementary results on unique ergodicity for such models with continuous state spaces.

As an implication of our analysis for controlled hidden Markov models, we provide new conditions for the existence of solutions to the average cost optimality equation for Partially Observable Markov Decision Processes, for which only limited results are available in the literature. We furthermore discuss implications on robustness to incorrect priors