More Risk-Sensitive Markov Decision Processes

Nicole Bäuerle

Institute of Stochastics, Karlsruhe Institute of Technology

We investigate the problem of minimizing a certainty equivalent of the total or discounted cost over a finite and an infinite horizon which is generated by a Markov Decision Process (MDP). The certainty equivalent is defined by $U^{-1}(EU(Y))$ where U is an increasing function. In contrast to a risk-neutral decision maker this optimization criterion takes the variability of the cost into account. It contains as a special case the classical risk-sensitive optimization criterion with an exponential utility. We show that this optimization problem can be solved by an ordinary MDP with extended state space and give conditions under which an optimal policy exists. In the case of an infinite time horizon we show that the minimal discounted cost can be obtained by value iteration and can be characterized as the unique solution of a fixed point equation using a 'sandwich' argument. Interestingly, it turns out that in case of a power utility, the problem simplifies and is of similar complexity than the exponential utility case, however has not been treated in the literature so far. We also establish the validity (and convergence) of the policy improvement method. A simple numerical example, namely the classical repeated casino game is considered to illustrate the influence of the certainty equivalent and its parameters.

The talk is based on joint work with U. Rieder.