Premium control with reinforcement learning

Lina Palmborg and Filip Lindskog

July 2022

Abstract

We consider a premium control problem in discrete time, inspired by [9] and [10], formulated in terms of a Markov decision process. In a simplified setting, the optimal premium rule can be derived with dynamic programming methods. However, these classical methods are not feasible in a more realistic setting due to the dimension of the state space. Hence, to combat the curse of dimensionality we explore reinforcement learning techniques, using linear function approximation. We illustrate the appropriateness of the approximate optimal premium rule compared with the true optimal premium rule in a simplified setting, and further demonstrate that the approximate optimal premium rule outperforms benchmark rules in a more realistic setting where classical approaches fail.