

Yao, Yi-Ching (Academia Sinica, Taiwan, R.O.C.)

Some Properties of Gittins Indices for Bernoulli and Normal Bandits

Abstract: For the classical multi-armed bandit problem with geometric discounting, the Gittins index plays a central role in that the optimal policy is to select at every stage the arm with the greatest Gittins index. In this talk, some properties of Gittins indices are derived and discussed for Bernoulli and normal bandits with (conjugate) beta and normal priors. A continuous-time version of the normal bandit problem is also considered where the Gittins index is determined by the solution of an optimal stopping problem involving Brownian motion. It is shown that this solution along with a continuity correction term provides a decent approximation for the original discrete-time normal bandit problem.