

Iterated Approximate Value Functions

B. O'Donoghue, Y. Wang, and S. Boyd

Proceedings European Control Conference, pages 3882-3888, Zurich, July 2013.

- [ECC paper](#)

In this paper we introduce a control policy which we refer to as the iterated approximate value function policy. The generation of this policy requires two stages, the first one carried out off-line, and the second stage carried out on-line. In the first stage we simultaneously compute a trajectory of moments of the state and action and a sequence of approximate value functions optimized to that trajectory. The next stage is to perform control using the generated sequence of approximate value functions. This gives a time-varying policy, even in the case where the optimal policy is time-invariant.

We restrict our attention to the case with linear dynamics and quadratically representable stage cost function. In this case the pre-computation stage requires the solution of a semidefinite program (SDP). Finding the control action at each time-period requires solving a small convex optimization problem which can be carried out quickly. We conclude with some examples.