

Approximate Dynamic Programming via Iterated Bellman Inequalities

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

International Journal of Robust and Nonlinear Control, 25(10):1472-1496, July 2015. Original manuscript posted June 2010.

- [adp_iter_bellman.pdf](#)

In this paper we introduce new methods for finding functions that lower bound the value function of a stochastic control problem, using an iterated form of the Bellman inequality. Our method is based on solving linear or semidefinite programs, and produces both a bound on the optimal objective, as well as a suboptimal policy that appears to work very well. These results extend and improve bounds obtained in a previous paper using a single Bellman inequality condition. We describe the methods in a general setting, and show how they can be applied in specific cases including the finite state case, constrained linear quadratic control, switched affine control, and multi-period portfolio investment.