

Shrinking-Horizon Dynamic Programming

J. Skaf, S. Boyd, and A. Zeevi

International Journal of Robust and Nonlinear Control, 20(17):1993-2002, 2010.

- [shdp.pdf](#)

We describe a heuristic control policy, for a general finite-horizon stochastic control problem, that can be used when the current process disturbance is not conditionally independent of previous disturbances, given the current state. At each time step, we approximate the distribution of future disturbances (conditioned on what has been observed) by a product distribution with the same marginals. We then carry out dynamic programming, using this modified future disturbance distribution, to find an optimal policy, and in particular, the optimal current action. We then execute only the optimal current action. At the next step we update the conditional distribution, and repeat the process, this time with a horizon reduced by one step. (This explains the name ‘shrinking horizon dynamic programming’.) We explain how the method can be thought of as an extension of model predictive control, and illustrate our method on two variations on a revenue management problem.

Page generated 2018-11-24 09:00:16 PST, by jemdoc.