The Curse of Dimensionality or the Blessing of Markovianity: Optimal Stopping as an Example

Robert J. Vanderbei*

April 9, 2014

Abstract

The beauty (and the blessing) of Markov processes is that they allow one to do computations on the *state space* E of the Markov process as opposed to working in the sample space of this process, which is at least as large as $E^{\mathbb{N}}$ and therefore huge relative to E.

Obviously, the tractability of a problem formulated on E depends on how big this space is. When E is a subset of a medium to high dimensional space, problems quickly become intractable. This fact is often referred to as the *curse of dimensionality*. But, even in high dimensions, E is small compared to $E^{\mathbb{N}}$. Hence, it is surprising that some recent papers have claimed to resolve the curse of dimensionality by working instead in the sample space. In this paper, we will review one such example.

Consider the *optimal stopping problem* associated with a discrete time Markov chain X_t on a finite (but large!) state space E. To avoid a trivial problem, assume that the chain has transient states (e.g., the chain could have one or more absorbing states). Let f denote a given non-negative payoff function. Assuming that the chain starts at a given point x_0, the problem, then, is to find a stopping time τ^* that

*Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA; e-mail: `rvdb@princeton.edu`

achieves the supremum over all stopping times τ of the expected value of f at X_τ:

$$v(x_0) = \sup_\tau \mathrm{E}_{x_0} f(X_\tau),$$

The classical method for solving this problem is to consider simultaneously all possible starting points x and then to note that, according to the *principle of dynamic programming*, v is the smallest function that satisfies

$$\begin{aligned} v(x) &\geq f(x) \qquad \text{for all } x && (1) \\ v(x) &\geq Pv(x) \qquad \text{for all } x, && (2) \end{aligned}$$

where P denotes the one-step transition matrix for the Markov process. In words, Eq. (1) says that v must majorize f and Eq. (2) says that v must be a P-*excessive* function.

Very few optimal stopping problems have simple closed-form solutions. In general, they require numerical methods to solve them. One approach is to compute v by the *method of successive approximation*:

$$\begin{aligned} v^{(0)}(x) &= 0 && \text{for all } x \\ v^{(1)}(x) &= \max\{f(x), Pv^{(0)}(x)\} && \text{for all } x \\ &\vdots \\ v^{(k+1)}(x) &= \max\{f(x), Pv^{(k)}(x)\} && \text{for all } x \\ &\vdots \end{aligned}$$

It can be shown that

$$v^{(k)}(x) = \sup_{\tau<k} \mathrm{E}_{x_0} f(X_\tau).$$

In words, it finds the optimal stopping time subject to the constraint that one must stop within k time steps. A second method for solving for $v(x_0)$ uses linear programming:

$$\begin{array}{lll} \text{minimize} & v(x_0) & \\ \text{subject to} & v(x) \geq f(x) & \text{for all } x \\ & v(x) \geq Pv(x) & \text{for all } x. \end{array}$$

These computational schemes are mathematically elegant but in practice they can be painfully slow if, say, the number of states is huge, which can be expected if the state space is a subset of a high-dimensional lattice. The fact that the numerical difficulty depends on the number of states is called the *curse of dimensionality*.

Davis and Karatzas (1994) claim to have resolved the curse of dimensionality in the following manner. Let M_t be an arbitrary mean zero martingale. Then,

$$\sup_\tau \mathrm{E}_{x_0} f(X_\tau) = \sup_\tau \mathrm{E}_{x_0} \left(f(X_\tau) - M_\tau\right) \leq \mathrm{E}_{x_0} \sup_t \left(f(X_t) - M_t\right).$$

Davis and Karatzas prove that the bound is tight and is achieved by the martingale part of the Doob/Meyer decomposition of Snell's envelope. In other words, we can solve either this problem

$$\sup_{\tau} \mathrm{E}_{x_0} f(X_\tau) = \inf_{M} \mathrm{E}_{x_0} \sup_{t} \left(f(X_t) - M_t \right) \tag{3}$$

or, if we know (or can guess) the optimal martingale, this one

$$\sup_{\tau} \mathrm{E}_{x_0} f(X_\tau) = \mathrm{E}_{x_0} \sup_{t} \left(f(X_t) - M_t^* \right). \tag{4}$$

Davis and Karatzas, and others after them, claim that it is easy to pick an almost correct martingale in many real-world examples and therefore that one can solve the optimal stopping problem by simply solving Eq. (4), which they claim is easier because it involves maximizing over a one-dimensional set, namely, time (we will ignore the fact that this set is infinite).

So, what is the optimal mean-zero martingale to subtract and how easy is it to guess it without knowing v? The martingale M^* is the martingale part of the Doob-Meyer decomposition of Snell's envelope. For optimal stopping, Snell's envelope is just $v(X_t)$ and the martingale part is simply

$$M_t^* = v(X_t) - v(X_0) - \sum_{s<t} Av(X_s),$$

where $A = P - I$ denotes the "infinitesimal" generator for X_t.

To check that the martingale part of the Snell envelope is correct, we compute as follows:

$$\begin{aligned} \sup_t (f(X_t) - M_t^*) &= \sup_t \left(f(X_t) - v(X_t) + v(X_0) + \sum_{s<t} Av(X_s) \right) \\ &= v(X_0) + \sup_t \left(f(X_t) - v(X_t) + \sum_{s<t} Av(X_s) \right) \\ &= v(X_0). \end{aligned}$$

The last equality follows from the observation that the supremum is just zero. To see that, we note that it is attained by setting t equal to the first hitting time of the set of states where $f = v$ (i.e., t is the optimal strategy in the original problem). At all prior times ($s < t$), $v = Pv$ and therefore $Av = 0$.

Rogers (2002) claims that, when pricing American options, a good guess for the martingale is the martingale part of the "payoff process", $f(X_t)$. In other words, he suggests using

$$M_t = f(X_t) - f(X_0) - \sum_{s<t} Af(X_s)$$

as a surrogate for M_t^*. With this choice, we get

$$\begin{aligned}\sup_t(f(X_t) - M_t) &= \sup_t\left(f(X_t) - f(X_t) + f(X_0) + \sum_{s<t} Af(X_s)\right) \\ &= f(X_0) + \sup_t \sum_{s<t} Af(X_s). \qquad (5)\end{aligned}$$

Conclusion

It seems that the only way to know the martingale M_t^* is to know the value function $v()$ and therefore the curse of dimensionality remains a curse. No?

1 An Example

Fix an integer n, put $\Delta x = 10^{-n}$ and let X_t be simple random walk on $[-1, 1]$ with step size Δx and with absorbing endpoints and let $f(x) = \cos(0.2 + 4x) + 1$. Figure 1 shows a plot of the payoff function f together with the value function v and an approximation to the value function computed using the expected value of the expression in (5) with the expected value computed using 1000 randomly generated trajectories of the random walk.

The approximate solution provides an upper bound on the value function, but at least for this example it is a terrible upper bound. Also, the approximate solution gives no clue as to the optimal stopping time τ^*, which is the first time that the random walk hits the set $\{x : v(x) = f(x)\}$.

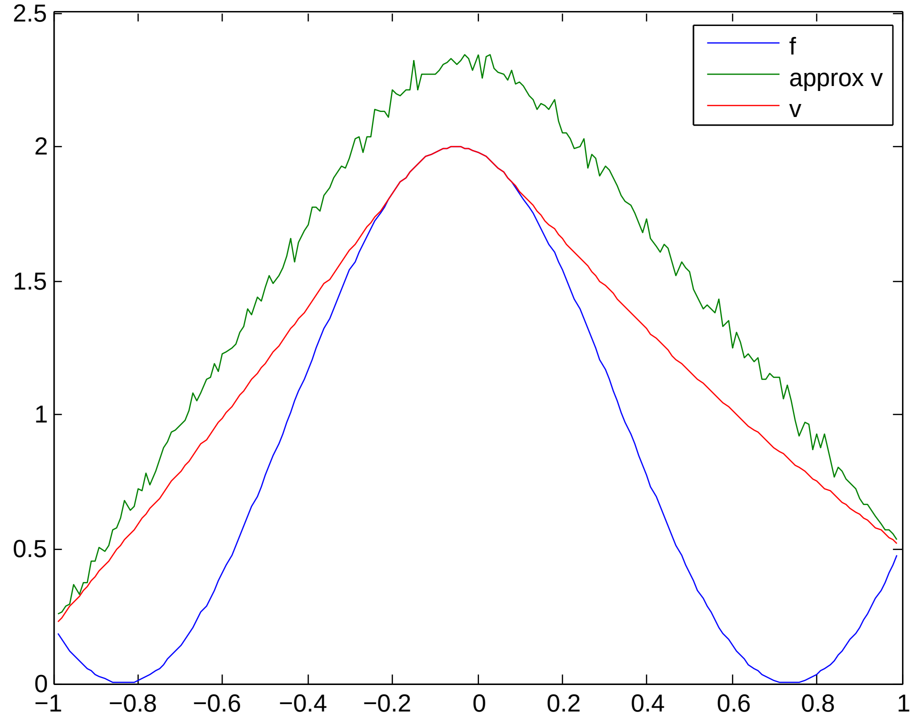

Figure 1: The exact and approximate value function associated with the payoff function shown.