

On Sequential Strategies for Loss Functions with Memory

Neri Merhav, Erik Ordentlich, Gadiel Seroussi, Marcelo J. Weinberger

ABSTRACT

The problem of optimal sequential decision for individual sequences, relative to a class of competing off-line reference strategies, is studied for general loss functions with memory. This problem is motivated by applications in which actions may have “long term” effects, or there is a cost for switching from one action to another. As a first step, we consider the case in which the reference strategies are taken from a finite set of generic “experts” We then focus on finite-state reference strategies, assuming finite action and observation spaces. We show that key properties that hold for finite-state strategies in the context of memoryless loss functions, do not carry over to the case of loss functions with memory. As a result, an infinite family of randomized finite-state strategies is seen to be the most appropriate reference class for this case, and the problem is basically different from its memoryless counterpart. Based on Vovk's exponential weighting technique, infinite-horizon on-line decision schemes are devised. For an arbitrary sequence of observations of length n , the excess normalized loss of these schemes relative to the best expert in a corresponding reference class is shown to be upper-bounded by an $O(n^{1/3})$ term in the case of a finite class, or an $O([\ln n]/n^{1/3})$ term for the class of randomized finite-state strategies. These results parallel the $O(n^{1/2})$ bounds attained by previous schemes for memoryless loss functions. By letting the number of states in the reference class grow, the notion of finite-state predictability is also extended.