
NICK HARVEY, Univeristy of British Columbia

How to make predictions using two experts, forever

A fundamental task in learning theory is to make repeated predictions given advice from experts. A famous algorithm for this task is the multiplicative weight method. It was independently invented in many areas including game theory, computational geometry, online learning, discrete optimization, and information theory.

As an algorithm designer, I am compelled to ask: is the multiplicative weight method the optimal algorithm for this problem? It turns out that it is the optimal algorithm under two assumptions: (1) the number of "experts" is large, and (2) the duration of the task is known in advance. We consider the optimality question after removing both of these assumptions. For the setting with just two experts in which the duration is unknown, we find the optimal algorithm. It is provably better than the multiplicative weight method.

The algorithm is derived in an unusual way. We reformulate the problem so that time progresses continuously rather than discretely. After doing so, the optimal algorithm can be found via a partial differential equation, which we solve analytically. This PDE solution is then used as the algorithm in discrete time. One would expect this approach to incur some discretization error, but bizarrely the discretization error is negative!

Joint work with Chris Liaw, Sikander Randhawa and Ed Perkins.