

## Defensive forecasting for decision making

Vladimir Vovk

Computer Learning Research Centre  
Department of Computer Science  
Royal Holloway, University of London  
Egham, Surrey TW20 0EX  
United Kingdom

We consider how to make probability forecasts of binary labels corresponding to vector objects. The basic observation is that for any continuous gambling strategy used for detecting disagreement between the forecasts and the actual labels, there exists a forecasting strategy whose forecasts are ideal as far as this gambling strategy is concerned. A forecasting strategy obtained in this way from a gambling strategy demonstrating a weak law of large numbers, which we call the K29 algorithm, is analyzed without making any assumptions about the way the data is generated. The algorithm is shown to be well calibrated and to have high resolution for big enough data sets and for a suitable choice of its parameter, a kernel on the Cartesian product of the forecast space  $[0, 1]$  and the object space. With a given utility function, good probability forecasts make it possible to make good decisions by maximizing expected utility. Applying this expected utility maximization principle to the forecasts produced by the K29 algorithm leads to surprising new results in prediction with expert advice, with the pool of experts comprising a functional space.