

The M5 Competition in Progress

Evangelos Spiliotis (National Technical University of Athens, spiliotis@fsu.gr)

Spyros Makridakis (Institute For the Future)

Vassilios Assimakopoulos (National Technical University of Athens)

The M5 competition follows on from the four previous M competitions, organized by Spyros Makridakis, whose purpose has been to advance the theory and practice of forecasting. The M5 differs from the previous four ones in five ways. First, it uses hierarchical unit sales data, generously made available by Walmart, starting at the product-store level and being aggregated to that of product departments, product categories, stores, and three geographical areas. Second, besides the time series data, it includes explanatory variables such as sell prices, promotions, days of the week, and special events that typically affect unit sales. Third, in addition to point forecasts, it assesses the distribution of uncertainty, as the participants are asked to provide information on nine indicative quantiles. Fourth, for the first time it focuses on series that display intermittency, i.e., sporadic demand including zeros. Finally, instead of a single competition, there are two parallel tracks, the first track concentrating on forecasting accuracy and the other on uncertainty. This session will present the preliminary results of both tracks of the M5 competition and discuss some initial conclusions from the analysis performed to the huge volumes of data being generated. The definite conclusions of the competition will be presented during the M5 Forecasting conference to be held in New York City on the 7th and 8th of December this year.