Forecasting Time Series with Multiple Levels of Seasonality

Abstract

Complex multiple seasonality is an important emerging challenge in time series forecasting. In this first part of this thesis, we propose models under a framework to forecast such time series. The framework segregates the task into two stages. In the first stage, the time series is aggregated and existing time series models such as regression, Box-Jenkins or TBATS, are used to fit this lower frequency data. In the second stage, additive or multiplicative seasonality at the higher frequency levels may be estimated using classical, or function-based methods. Finally, the estimates from the two stages are combined. The proposed framework is capable of handling covariates, is low on computational complexity and does not suffer from instability in forecasting (for example overfitting). We compare the forecast accuracy of this framework with a popular method proposed by De Livera et al. [2011] and show that the proposed framework performs favourably in both short (day ahead forecast) and long term forecast horizons (one year ahead forecast) in a few case studies.

In the second part of the thesis we describe the development of an `R` package for the framework in order to make it more accessible to practitioners and researchers.

In the third part of the thesis, we look at the impact of data frequency on forecast accuracy. In order to forecast a variable at a particular frequency (say hourly frequency), obtaining the data at higher frequency (for example, fifteen minute frequency) may be costly in terms of transmission (power and bandwidth) and storage (memory). However, if the data is collected only at the desired frequency, which is then used for prediction, we might be losing out on vital information useful for a better prediction. We first establish that high frequency data is always

useful, but then show that there are situations when a compromise may be made between the desired accuracy and the higher memory, power and bandwidth needed to process high frequency data. We introduce Wiener-Kolmogorov filtering for a stationary time series with a single level of seasonality and derive mean squared error expressions for forecasting data at high frequency and for forecasting aggregated data at the desired frequency. These provide the information needed to decide when the compromise can be safely made. We then generalize the result to multiple levels of seasonality. We verify our results using both simulated data and two datasets.

In the last part of the thesis, we examine optimal combinations of individual forecasts. Forecast combining is a well-known method of improving forecast accuracy. The best linear combination in terms of minimising the mean squared error (MSE) of the combination is one of the ways to combine forecasts. If this optimal combination entails negative weights for some forecasts, it may not be very appealing from the perspective of implementation and interpretation. In particular, while combining two forecasts, the optimal weights turn out to be positive when the two forecasts are negatively correlated. However, if they are positively correlated, the optimal weights are positive only if the correlation is less than a threshold value. In this work, we obtain a necessary and sufficient condition for the optimum linear combination of an arbitrary number of good forecasts to be a convex combination. We then observe that the condition for convexity, when there are more than three good forecasts to be combined, does not have a simple interpretation. We propose checking for pair-wise convexity in such cases.