Topics of the lab: State space models based on exponential smoothing and automated model fitting with forecast package

At the beginning of the course we considered 4 time series forecasting methods which are based on exponential smoothing: simple exponential smoothing, Holt's method, Holt-Winters methods. There are actually many more such methods, which are systematically discussed in the book "Rob J. Hyndman, Anne B Koehler, J. Keith Ord, Ralph D. Snyder, Forecasting with Exponential smoothing: The State Space Approach." All those methods are implemented b R. Hyndman in the R package forecast.

The basis of the models is that the time series is formed by combining together three components T (trend), S (seasonality) and E(errors). Each component can be added to previous ones or multiplied by previous components, seasonality can also be missing. Three different types of trend are considered - the trend curve for time t is described by the level l_t, change rate parameter b_t and for some trend types also by a damping factor ϕ. The 5 possible trend curves are as follows:

1. No trend: $T_{t+h|t} = l_t$
2. Additive: $T_{t+h|t} = l_t + b_t h$
3. Damped additive: $T_{t+h|t} = l_t + (\sum_{i=1}^{h} \phi^i) b_t$, $0 < \phi < 1$
4. Multiplicative: $T_{t+h|t} = l_t b_t^h$
5. Damped multiplicative: $T_{t+h|t} = l_t b_t^{\sum_{i=1}^{h} \phi^i}$, $0 < \phi < 1$.

So all models can be described by a triple of descriptors, where the first one describes the error type (additive A or multiplicative M), the second one the main type of trend (no trend N, additive A or multiplicative M) and the third one the type of seasonality (N,A or M). In addition to the main type one should indicate weather the normal or damped version of the trend is used. All such models can be considered to be examples of state space models, where the state consists of the level l_t, the change rate b_t and seasonality parameters $S_t, S_{t-1}, \ldots, S_{t-p+1}$, where p denotes the length of the period and the changes of state parameters is based on exponential smoothing ideas (smoothing parameters are α, β and γ).

There are also programs which try to find the best ARIMA model by trying to fit many models and choosing the best one according to some criterion. One such tool is auto.arima() in the forecast package. As it is true about using any complicated tool, one should know how to use command correctly - which options should be modified in which cases, what are the limitations of the command and so on. In the lab we'll discuss some aspects of using the command.

Exercises:

1. Install the package forecast. View the help for the command ets() and ask if something remains unclear. Find the best exponential smoothing models for the time series considered in the first lab. Are the resulting models suitable for predictions?
2. Read the help for auto.arima(). Try automatic fitting of ARIMA models to the time series time series in the data set `lab15.csv` and to the series we considered in the first lab. Are the models the same you would find by the manual procedures considered in this course?