

Topics of the lab: fitting AR, MA and ARMA models to stationary time series, checking the goodness of fit.

Our procedure of model selection is as follows:

1. Look at the graphs of a time series, its autocorrelations and partial autocorrelations to decide if the series is stationary and to determine the type of a model (and the number of its parameters) which is a good starting point for modelling the series. Some facts about ARMA(p,q) models which are useful in model selection:
 - Theoretical partial autocorrelations of AR(p) models (or, equivalently ARMA(p,0) models) are equal to 0 starting from the lag $p+1$. Additional fact about AR(1) models is that the autocorrelations should behave like w^k (where $|w| < 1$) (thus ratios of consecutive autocorrelations should be equal to a constant).
 - Theoretical autocorrelations of MA(q) models are equal to 0 starting from the lag $q+1$. Additionally, the partial autocorrelations of MA(1) models should behave like w^k (where $|w| < 1$).
 - Theoretical autocorrelations of ARMA(1,q) models behave like $\rho_{k+1} = w \cdot \rho_k$ (where $|w| < 1$) starting from $k = q$). In the case of ARMA(1,1) models both theoretical autocorrelations and partial autocorrelations are decreasing towards 0 after lag 1 like being multiplied by a constant whose absolute value is less than 1.
2. Fit the selected model to the time series (let the software to find the best values for the parameters).
3. Check the suitability of the fitted model. The main criteria is the randomness of the one-step prediction errors, which we check by looking at autocorrelations of the residuals of the model and the p-value of Ljung-Box statistic in the case of a suitable group (10-20 in the case of non-seasonal series, 24-36 in the case of monthly data) of autocorrelations of the prediction errors. If the autocorrelations of residuals do not correspond to uncorrelated series, we try to fit a better model.

Exercises of the lab:

1. Find suitable models to all those time series in `lab6.csv` and `lab6_2.dat`, which can be considered to be stationary. For fitting models we are going to use the command `arima`. To specify the model to be fitted the command has a parameter `order`, which should be a vector of three integers of the form `c(p,0,q)` for ARMA(p,q) models. Suitability of a fitted model can be checked by the command `tsdiag()`, which shows the graphs of the residuals, their autocorrelations and also p-values of Ljung-Box statistics in the case of various group sizes. Although the p-values are not correct (the number of parameters in the fitted model is not taken into account), the effect of the number of parameters on the p-value is quite small for larger group sizes, so you can make decisions about suitability of a model based on the results of the `tsdiag()` command (provided the parameter `gof.lag` is given a large enough value).
2. In addition to being able to fit a good model to data, it is also important to understand, how the formula for the best model looks like. Write down on a paper the formulas of models you fitted to the time series data.

Remark. In the case of using `arima()` command, the coefficients of the fitted models are given exactly as they appear on the right hand side of the model formula and the term `intercept` is an estimate of the overall mean value of the series.