Time Series Analysis

Week 13 –

Peder Bacher
Department of Applied Mathematics and Computer Science
Technical University of Denmark

May 3, 2024

The tool box

- General linear model
  - Global trend models
  - Local trend models
- Univariate linear dynamic models
  - ARIMA
  - ARMAX
  - Transfer function models
- MARIMA
- State-space models
- Recursive/adaptive linear regression models

The tool box - Global trend models

General form:

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{\theta} + \boldsymbol{\epsilon}$$

Example:

$$Y_t = \begin{bmatrix} 1 & t & t^2 & u_t \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \\ \theta_4 \end{bmatrix} + \epsilon_t$$

Decisions:

- What do you want to regress on?
- OLS or WLS? And how to decide weights for WLS.

Notes:

- Can be used for repeated experiments.

The tool box - Global trend models

General form:

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{\theta} + \boldsymbol{\epsilon}$$

Example:

$$Y_t = \begin{bmatrix} 1 & t & t^2 & u_t \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \\ \theta_4 \end{bmatrix} + \epsilon_t$$

Decisions:

- What do you want to regress on?
- OLS or WLS? And how to decide weights for WLS.

Notes:

- Can be used for repeated experiments.
- Examples: Growth of humans, annual patterns.

The tool box - Local Trend Models

General form:

$$Y_{t+n|t} = f(n)\boldsymbol{\theta}_t + \epsilon_{t+n}$$

Example:

$$Y_{t+n|t} = \begin{bmatrix} 1 & n & n^2 \end{bmatrix} \begin{bmatrix} \mu_t \\ \alpha_t \\ \beta_t \end{bmatrix} + \epsilon_{t+n}$$

Decisions:

- Order of trend.
- Memory - should be tuned to horizon of interest.

Notes:

- Can pretty much always be used for short-term predictions - Fall-back method.
- Can be interpreted as exponential smoothing with a trend - Robust baseline for comparisons.
- Can not incorporate any domain knowledge.

The tool box - Local Trend Models

General form:

$$Y_{t+n|t} = f(n)\boldsymbol{\theta}_t + \epsilon_{t+n}$$

Example:

$$Y_{t+n|t} = \begin{bmatrix} 1 & n & n^2 \end{bmatrix} \begin{bmatrix} \mu_t \\ \alpha_t \\ \beta_t \end{bmatrix} + \epsilon_{t+n}$$

Decisions:

- Order of trend.
- Memory - should be tuned to horizon of interest.

Notes:

- Can pretty much always be used for short-term predictions - Fall-back method.
- Can be interpreted as exponential smoothing with a trend - Robust baseline for comparisons.
- Can not incorporate any domain knowledge.
- Examples: Sales, financial assets, sensor alerts, de-trending.

The tool box - Univariate linear dynamic models - (Seasonal) ARIMA

General form:

$$\phi(B)\Phi(B^s)\nabla^d\nabla_s^D Y_t = \theta(B)\Theta(B^s)\epsilon_t$$

Decisions:

- Seasonality.
- Differencing.
- Order of AR and MA components.

Assumptions:

- No measurement noise.
- Normal distributed noise.

Notes:

- No inputs.
- Implicit modelling of dynamics and relations.
- MA-part can be interpreted as a model for coloured noise.

The tool box - Univariate linear dynamic models - (Seasonal) ARIMA

General form:

$$\phi(B)\Phi(B^s)\nabla^d\nabla_s^D Y_t = \theta(B)\Theta(B^s)\epsilon_t$$

Decisions:

- Seasonality.
- Differencing.
- Order of AR and MA components.

Assumptions:

- No measurement noise.
- Normal distributed noise.

Notes:

- No inputs.
- Implicit modelling of dynamics and relations.
- MA-part can be interpreted as a model for coloured noise.
- Examples: House Prices, sales, some physical systems.

The tool box - Univariate linear dynamic models - ARMAX

General form:

$$\varphi(B)\, Y_t = \omega(B) B^b X_t + \theta(B)\epsilon_t$$

Decisions:

- Order of AR and MA components.
- Regression on X_t, i.e. ω and b.

Assumptions:

- No measurement noise.
- Normal distributed noise.

Notes:

- Used to extend ARMA models to include regressors
- Pre-whitening is needed to find ccf between Y and X.
- Implicit modelling of dynamics and relations.
- MA-part can be interpreted as a model for coloured noise.

The tool box - Univariate linear dynamic models - ARMAX

General form:

$$\varphi(B)Y_t = \omega(B)B^b X_t + \theta(B)\epsilon_t$$

Decisions:

- Order of AR and MA components.
- Regression on X_t, i.e. ω and b.

Assumptions:

- No measurement noise.
- Normal distributed noise.

Notes:

- Used to extend ARMA models to include regressors
- Pre-whitening is needed to find ccf between Y and X.
- Implicit modelling of dynamics and relations.
- MA-part can be interpreted as a model for coloured noise.
- Examples: House prices, sales, physical systems under control.

The tool box - Univariate linear dynamic models - Transfer Function Models

General form:

$$Y_t = \frac{\omega(B)}{\delta(B)} B^b X_t + \frac{\theta(B)}{\varphi(B)} \epsilon_t$$

Decisions:

- Order of the four polynomials

Assumptions:

- Normal distributed noise.

Notes:

- Implicit modelling of dynamics and relations.
- Simplifies to ARMAX for $\delta(B) = \varphi(B)$.
- Most often used for frequency-domain identification and control.

The tool box - Univariate linear dynamic models - Transfer Function Models

General form:

$$Y_t = \frac{\omega(B)}{\delta(B)} B^b X_t + \frac{\theta(B)}{\varphi(B)} \epsilon_t$$

Decisions:

- Order of the four polynomials

Assumptions:

- Normal distributed noise.

Notes:

- Implicit modelling of dynamics and relations.
- Simplifies to ARMAX for $\delta(B) = \varphi(B)$.
- Most often used for frequency-domain identification and control.
- Examples: Mechanical systems, Race car suspension systems.

The tool box - MARIMA

General form:

$$\boldsymbol{\phi}(B)\,\boldsymbol{Y}_t = \boldsymbol{\omega}(B)B^b\boldsymbol{X}_t + \boldsymbol{\theta}(B)\boldsymbol{\epsilon}_t,$$

where $\boldsymbol{\phi}$ and $\boldsymbol{\theta}$ are matrices, and $\boldsymbol{Y}_t$, $\boldsymbol{X}_t$ and $\boldsymbol{\epsilon}_t$ are vectors.

Decisions:

- Order of the AR and MA components for each variable.
- Regression model for each input.

Assumptions:

- Normal distributed noise.
- No measurement noise.

Notes:

- Used for modelling time series that effect each other.
- Only needed when causality goes both ways.
- Implicit modelling of dynamics and relations.

The tool box - MARIMA

General form:

$$\boldsymbol{\phi}(B)\,\boldsymbol{Y}_t = \boldsymbol{\omega}(B)B^b\boldsymbol{X}_t + \boldsymbol{\theta}(B)\boldsymbol{\epsilon}_t,$$

where $\boldsymbol{\phi}$ and $\boldsymbol{\theta}$ are matrices, and $\boldsymbol{Y}_t$, $\boldsymbol{X}_t$ and $\boldsymbol{\epsilon}_t$ are vectors.

Decisions:

- Order of the AR and MA components for each variable.
- Regression model for each input.

Assumptions:

- Normal distributed noise.
- No measurement noise.

Notes:

- Used for modelling time series that effect each other.
- Only needed when causality goes both ways.
- Implicit modelling of dynamics and relations.
- Examples: House prices, homicide/suicide rates, political party popularity.

The tool box - Linear State-space models

General form:

$$\boldsymbol{X}_{t+1} = \boldsymbol{A}\boldsymbol{X}_t + \boldsymbol{B}\boldsymbol{U}_t + \boldsymbol{\epsilon}_t$$
$$\boldsymbol{Y}_t = \boldsymbol{C}\boldsymbol{X}_t + \boldsymbol{e}_t$$

Decisions:

- Structure of $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$.

Assumptions:

- Normal distributed noise.
- White noise.

Notes:

- Used for systems that are subject to both measurement and process noise.
- Used for systems with hidden (i.e. unobserved) states.
- This usually applies to physical systems.
- Parameters can be interpreted as hidden states.
- Handles missing observations seamlessly.

The tool box - Linear State-space models

General form:

$$\boldsymbol{X}_{t+1} = \boldsymbol{A}\boldsymbol{X}_t + \boldsymbol{B}\boldsymbol{U}_t + \boldsymbol{\epsilon}_t$$
$$\boldsymbol{Y}_t = \boldsymbol{C}\boldsymbol{X}_t + \boldsymbol{e}_t$$

Decisions:

- Structure of $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$.

Assumptions:

- Normal distributed noise.
- White noise.

Notes:

- Used for systems that are subject to both measurement and process noise.
- Used for systems with hidden (i.e. unobserved) states.
- This usually applies to physical systems.
- Parameters can be interpreted as hidden states.
- Handles missing observations seamlessly.
- Examples: Temperatures in buildings, blood sugar/insulin levels in humans, position of satellites.

The tool box - Recursive (Pseudo) Linear Regression

General form:

$$Y_{t+1|t} = x_t \boldsymbol{\theta}_t + \epsilon_{t+1}, \qquad \text{or sometimes} \qquad Y_{t+1|t} = x_t(\theta)\boldsymbol{\theta}_t + \epsilon_{t+1}.$$

Decisions:

- Regression model, i.e. variables in x_t.
- Memory, i.e. λ - should be tuned to horizon of interest.

Assumptions:

- Normal distributed noise.
- White noise.

Notes:

- Used when we want to repeatedly update our parameter estimates.
- Used when we believe that parameters are changing over time.
- Used when we know that our model is too simple.
- Variable forgetting can be used for models that we believe are not changing most of the time.
- Can be applied to ARI(MA) models using (pseudo)-linear recursive regression.

The tool box - Recursive (Pseudo) Linear Regression

General form:

$$Y_{t+1|t} = x_t \boldsymbol{\theta}_t + \epsilon_{t+1}, \qquad \text{or sometimes} \qquad Y_{t+1|t} = x_t(\theta) \boldsymbol{\theta}_t + \epsilon_{t+1}.$$

Decisions:

- Regression model, i.e. variables in x_t.
- Memory, i.e. λ - should be tuned to horizon of interest.

Assumptions:

- Normal distributed noise.
- White noise.

Notes:

- Used when we want to repeatedly update our parameter estimates.
- Used when we believe that parameters are changing over time.
- Used when we know that our model is too simple.
- Variable forgetting can be used for models that we believe are not changing most of the time.
- Can be applied to ARI(MA) models using (pseudo)-linear recursive regression.
- Examples: District heating consumption, pretty much everything that the other models can be used for.

Courses where 02417 is really used

- 02427 Advanced Time Series Analysis
- 02407 Stochastic Processes - Probability 2
- 02425 Diffusions and Stochastic Differential Equations
- 02421 Stochastic Adaptive Control

THE END