An introduction to time series models

ITT9
January 28th, 2019

Matt Nunes, University of Bath

Time series analysis

Time series analysis simply refers to the analysis of data collected / indexed over time. Such data is observed in a wide range of scientific areas of interest, e.g. industrial process monitoring, climate modelling, official statistics.

In particular, our aim is to build **realistic models** of such data which account for possible complex **temporal dependencies**.

Time series analysis

Time series analysis simply refers to the analysis of data collected / indexed over time. Such data is observed in a wide range of scientific areas of interest, e.g. industrial process monitoring, climate modelling, official statistics.

In particular, our aim is to build **realistic models** of such data which account for possible complex **temporal dependencies**.

Analysis tasks after modelling include

- forecasting (prediction)
- classification / distinguishing series
- detection of changes, identifying patterns or periodicities etc.

Time series analysis

Time series analysis simply refers to the analysis of data collected / indexed over time. Such data is observed in a wide range of scientific areas of interest, e.g. industrial process monitoring, climate modelling, official statistics.

In particular, our aim is to build **realistic models** of such data which account for possible complex **temporal dependencies**.

Analysis tasks after modelling include

- forecasting (prediction)
- classification / distinguishing series
- detection of changes, identifying patterns or periodicities etc.

Notation:

- A (real-valued, stationary) time series will be denoted by $\{X_t\}_{t \in \mathbb{Z}}$, with a corresponding realisation of X_t being x_t.

Stationarity

In order to do inference, it is often assumed some sort of invariance of time series, i.e. the statistical characteristics of the series do not change over time (**stationarity**).

Types of stationarity:

- **First order**: The mean of the time series is the same over time
- **Second order / covariance / weak stationarity**: If the mean is constant for all t and if for any t and k, $\gamma_X(h) = \mathrm{cov}(X_t, X_{t+k})$ only depends on the lag difference k.

Stationarity

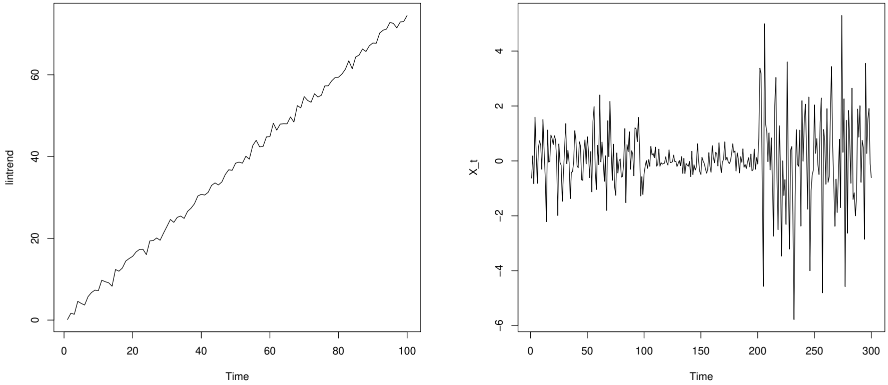

Figure: Types of (non)stationarity: linear trend (left); non-constant variance (right).

Some popular time series models: AR(p)

Motivation: Recall from linear regression, we predict a response Y given some covariates X_j, so we model Y_i as

$$Y_i = \sum_{j=1}^{p} a_j X_{ij} + \varepsilon_i,$$

with $\mathbb{E}(\varepsilon_i | X_{ij}) = 0$ and typically ε_i and X_{ij} independent.

For time series, we can similarly predict a future observation from the current and past observations

$$X_t = \sum_{j=1}^{p} a_j X_{t-j} + \varepsilon_t.$$

This is the **autoregressive model (of order p)**.

Some popular time series models: MA(q)

Let $\{X_t\}$ be a time series. We say X_t has a **moving average of order q** (MA(q) for short) representation if

$$X_t = \sum_{j=0}^{q} \psi_j \varepsilon_{t-j},$$

where $\{\varepsilon_t\}$ are IID random variables with zero mean and finite variance (i.e. white noise).

In other words, the series is modelled as a linear combination of the previous noise.

Some popular time series models: MA(q)

Let $\{X_t\}$ be a time series. We say X_t has a **moving average of order q** (MA(q) for short) representation if

$$X_t = \sum_{j=0}^{q} \psi_j \varepsilon_{t-j},$$

where $\{\varepsilon_t\}$ are IID random variables with zero mean and finite variance (i.e. white noise).

In other words, the series is modelled as a linear combination of the previous noise.

We can combine autoregressive and moving average models to form ARMA models.

ARMA processes: model selection

- Looking at the autocorrelation function (ACF) and partial autocorrelation function can give an idea about how to choose model AR and MA orders (look for where the plots “cut off”)

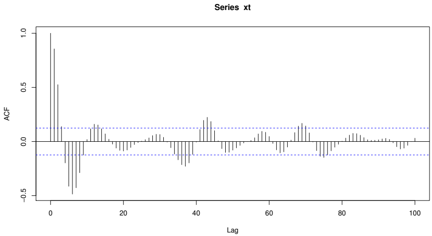

Figure: ACF and PACF of an AR(2) process; notice the characteristic “cut off” and damped exponential pattern of the plots.

- we can also use more formal model selection procedures like the AIC.

Integrated models for (trend) nonstationarity

Now suppose X_t has a polynomial trend behaviour ($\mu_t = \sum_{j=0}^{k} \beta_j t^j$).

Then if we difference the series k times (repeatedly take differences of successive observations), then the resulting series will be stationary.

This leads to the **integrated ARMA model**: Define $\nabla x_t = x_t - x_{t-1}$, then a process x_t is said to be ARIMA(p,d,q) if $\nabla^d x_t$ is ARMA(p,q).

Integrated models for (trend) nonstationarity

Example: Stationary process with a linear trend:

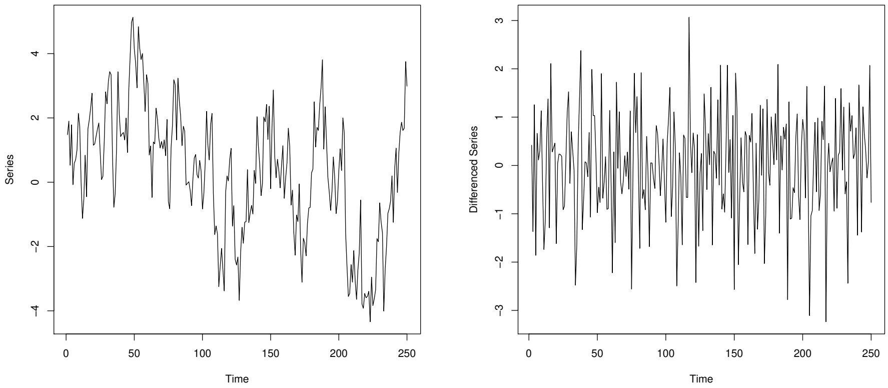

Figure: Effect of differencing: original series (left); differenced series (right).

Modelling seasonality

We can also extend the models we've seen to seasonal components, in a similar manner to integrated models.

Suppose a seasonal cycle lasts for s timepoints, i.e. the behaviour of the series is similar at a lag of s. Then if we difference the series **at lag s**,

$$y_t = \nabla_s x_t = x_t - x_{t-s},$$

this will remove the seasonality.

Modelling seasonality

We can also extend the models we've seen to seasonal components, in a similar manner to integrated models.

Suppose a seasonal cycle lasts for s timepoints, i.e. the behaviour of the series is similar at a lag of s. Then if we difference the series **at lag s**,

$$y_t = \nabla_s x_t = x_t - x_{t-s},$$

this will remove the seasonality.

More generally, we can have a ARIMA for the seasonal part of a time series $(P, D, Q)_s$, with the same interpretation of the orders.

Modelling seasonality

We can also extend the models we've seen to seasonal components, in a similar manner to integrated models.

Suppose a seasonal cycle lasts for s timepoints, i.e. the behaviour of the series is similar at a lag of s. Then if we difference the series **at lag s**,

$$y_t = \nabla_s x_t = x_t - x_{t-s},$$

this will remove the seasonality.

More generally, we can have a ARIMA for the seasonal part of a time series $(P, D, Q)_s$, with the same interpretation of the orders.

Putting this together a flexible model is the **SARIMA model**: $ARIMA(p, d, q) \times ARIMA(P, D, Q)_s$.

This allows for nonseasonal and seasonal components.

Modelling seasonality

Example: `co2` time series representing monthly CO_2.

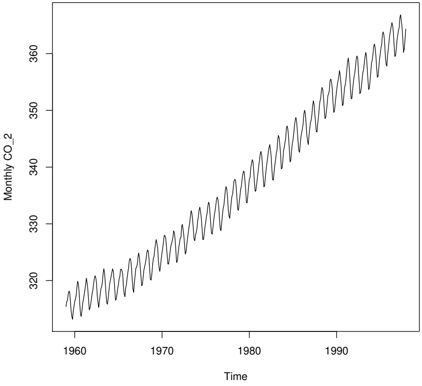

Figure: Time series featuring trend and yearly seasonality.

Spectral analysis: frequency domain representations

In many applications, time series will exhibit periodicities or oscillations, which may occur at differing rates.

These periodicities may be difficult to discern in the time domain.

Spectral / frequency domain analysis aims to capture these features, and provide extra insight and properties of data.

Spectral analysis: frequency domain representations

In many applications, time series will exhibit periodicities or oscillations, which may occur at differing rates.

These periodicities may be difficult to discern in the time domain.

Spectral / frequency domain analysis aims to capture these features, and provide extra insight and properties of data.

Main idea:

decompose a (stationary) series in terms of sinusoids at different frequencies ω_j with random, uncorrelated amplitudes.

Spectral analysis: frequency domain representations

Suppose $X_t = \sum_{j=1}^{k} A_j \sin(2\pi\omega_j t) + B_j \cos(2\pi\omega_j t)$, with A, B uncorrelated, mean zero, with variance σ_j^2 (mixture of sinusoids at different frequencies and amplitudes).

Then,

$$\gamma(h) = \sum_{j=1}^{k} \sigma_j^2 \cos(2\pi\omega_j h).$$

(This follows from the uncorrelatedness of A_j and B_j).

Spectral analysis: frequency domain representations

Suppose $X_t = \sum_{j=1}^{k} A_j \sin(2\pi\omega_j t) + B_j \cos(2\pi\omega_j t)$, with A, B uncorrelated, mean zero, with variance σ_j^2 (mixture of sinusoids at different frequencies and amplitudes).
Then,

$$\gamma(h) = \sum_{j=1}^{k} \sigma_j^2 \cos(2\pi\omega_j h).$$

(This follows from the uncorrelatedness of A_j and B_j). In particular, setting $h = 0$, we have

$$\text{var}(X_t) = \gamma(0) = \sum_{j=1}^{k} \sigma_j^2.$$

In other words, we can decompose the autocovariance / variance of the process via the sinusoidal components of the series X_t (via the Fourier transform).

The spectral density: some comments

Some comments of the spectral density:

- Interpretation: A stationary time series can be (approximately) expressed as a random linear combination of sines and cosines at different frequencies).
- The spectral density is positive
- The spectral density contains **the same information** as the autocovariance, just expressed differently (cf. Parseval's theorem).
- The spectral density is even and periodic (hence we can restrict our attention to e.g. $\omega \in (0, 1/2)$).

Spectrum example

Let $X_t = 2\cos(2\pi 6t/100) + 4\cos(2\pi 10t/100) + 6\cos(2\pi 40t/100)$.

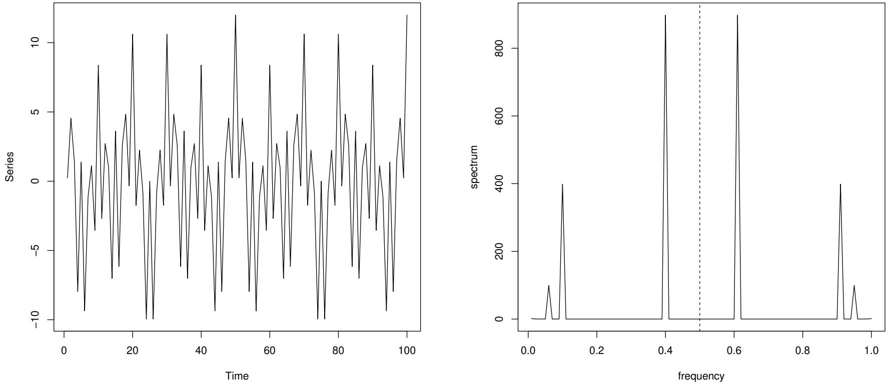

Figure: Spectrum estimate of X_t, featuring three periodicities at distinct frequencies ("full" frequency range).

Spectrum example

Let $X_t = 2\cos(2\pi 6t/100) + 4\cos(2\pi 10t/100) + 6\cos(2\pi 40t/100)$.

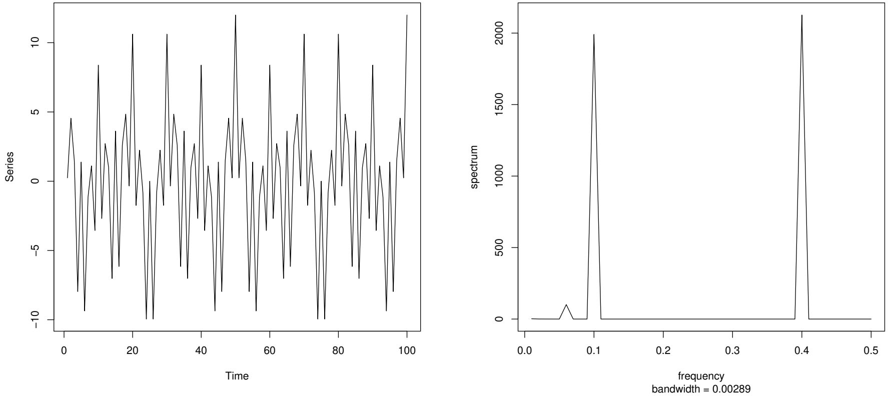

Figure: Spectrum estimate of X_t, featuring three periodicities at distinct frequencies.

Periodogram examples

Let X_t be the `soi` (Southern Oscillation Index) series (below).

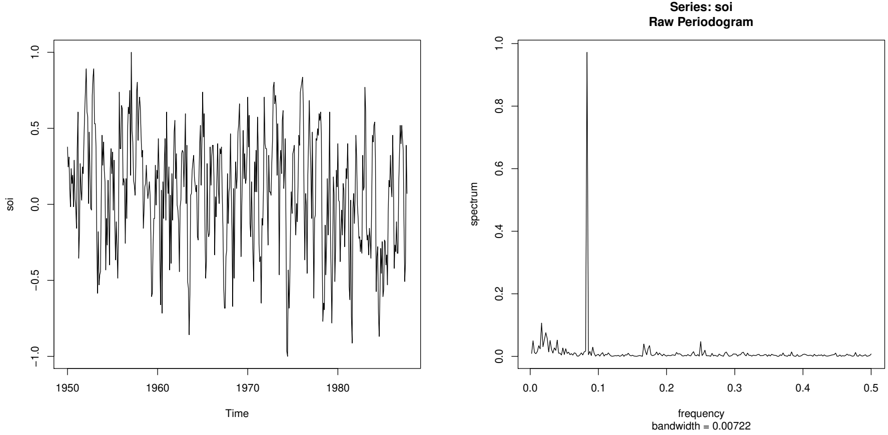

Figure: Spectrum estimate of the Southern Oscillation Index.

Example analysis tasks

- Forecasting:
  - There are many ways to forecast a time series, depending on your intuition and the model. For example, one could use: a naive estimator, a moving average, exponential smoothing
  - if there are trend and seasonal components, these can also be taken into account by using similar procedures, or using the model form
- Changepoint detection:
  - What kind of changes are you looking for? Abrupt or ramping?
  - are you looking for changes in mean, variance or autocovariance? Different methods for each

Forecasting example

Example: `co2` time series representing monthly CO_2.

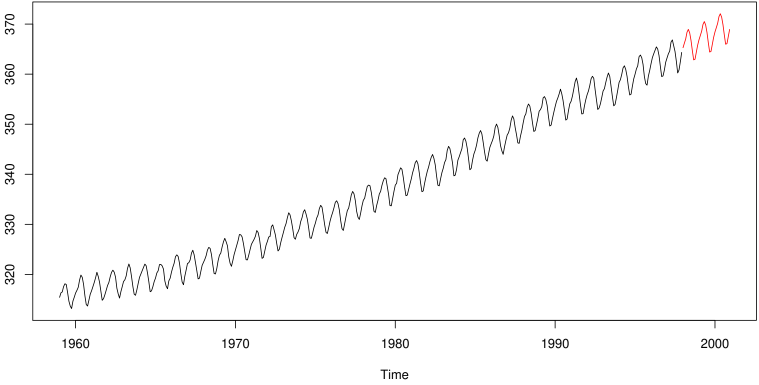

Figure: Original series, together with forecasted values.

Changing structure

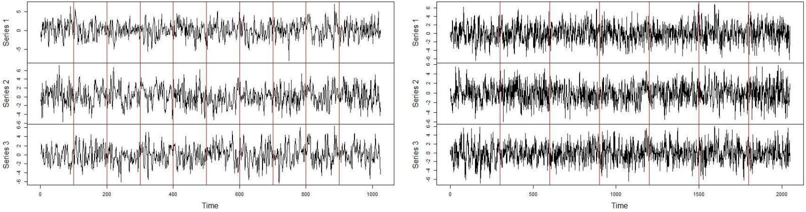

Figure: Two similar-looking series with different underlying behaviour (environmental sensor data).

Box-Jenkins modelling time series

Box / Jenkins suggest the following procedure for time series:

1. Examine the time series and remove any seasonality
2. Remove any trend by appropriate differencing (usually ∇_1 or ∇_2)
3. Fit appropriate ARMA model
4. Begin forecasting!

Other Remarks and considerations

There are many other issues in time series which are relevant, including

- multivariate models: including dependence **between** series; non-Gaussian errors: transform (e.g. via log) or use a count process model
- addition of covariates (exogenous variables): what's driving the behaviour?
- second order nonstationarity: e.g locally stationary models
- Dynamically changing situations: dlms
- In `R`, see the `forecast`, `changepoint`, `dlm` packages (in particular `auto.arima`, `stl`, `adf.test` functions)

References

Chatfield, C. (2004) *The Analysis of Time Series: An Introduction. Chapman & Hall/CRC Pr I Llc.*

Hyndman, R. J. and Athanasopoulos, G. (2018) Forecasting: Principles and Practice. OTexts.

*Killick, R. and Eckley, I. A. (2014) changepoint: An r package for changepoint analysis. Journal of Statistical Software, **58**, 1–19.*

Koopmans, L. H. (1995) The Spectral Analysis of Time Series. Academic Press.

Petris, G., Petrone, S. and Campagnole, P. (2009) Dynamic linear models with R. Springer.

Shumway, R. H. and Stoffer, D. S. (2017) Time Series Analysis and Its Applications. With R Examples. Springer.

Challenges in time series analysis

- Modelling and differentiating series with complex dependencies dependence
- Online analysis of fast-changing time series / data streams
- Estimation in high dimensions computationally intensive
- Tools for “mixed rate” sampling series