

Day 3: Self-guided practical - finding the best time series model for your data set

Andrew Parnell

Introduction

By now you've seen pretty much everything that the `forecast` package has to offer, and a little bit of the Bayesian models too. This should give you just enough information to get started on your own data set. However, if you still haven't seen anything that looks like your data yet, try analysing one of the example data sets in the `data` folder, or go online and find a good data set. There are lots of good sites, including:

- The R `datasets` library. Just type `library(help = "datasets")`
- The [Time Series Data Library](#)
- Some of the forecasting competition sites, such as [Kaggle](#), [CrowdAnalytix](#), and [TunEdit](#)

Tasks

You know the drill by now! Take your data set and:

1. Plot it
2. Transform it if necessary (by e.g. differencing or Box-Cox)
3. Fit an ARIMA, ARIMAX, or similar model
4. Use a suitable test set, AIC, or rolling origin forecast to decide on the best models
5. Check the residuals
6. Predict the future!

If you get stuck please ask for help.