

Day 4: Self-guided practical - start analysing your own data set with Bayes

Andrew Parnell

Introduction

If your time series is in any way irregular; repeated observations, missing values, multivariate, etc, you are not going to be able to analyse it fully with the `forecast` package. I would strongly recommend moving over to using JAGS or Stan to fit your models. In this practical you will move over from using `forecast` to fitting a fully Bayesian statistical model.

Tasks

1. Take your data set (or the one you were using yesterday)
2. Pick either JAGS or Stan (if you choose Stan you might need to install the package). Go back and revise one of the simpler models that reasonably matches your data (e.g. linear regression, AR(1), etc) and check you can get it to run on example data
3. Keep the model the same, but change the setup code so your data set works with it
4. Start to expand the model. This might involve:
 - Adding in explanatory variables
 - Adding in extra AR or MA components
 - Including Fourier or seasonal parts
 - Including stochastic volatility or heteroskedastic components
5. When you have got a model running, make sure to check convergence using the R-hat values. If the model hasn't converged try increasing the number of iterations. If it still doesn't converge, then it might be a problem with the model definition. Get some help
6. Try a few different similar models and compare using the DIC if using JAGS or WAIC if using Stan.
7. See if you can create some forecasts (if that is your goal) or interpret your model output