Time Series Laboratory

Computing in Weber Classrooms 205-206:

- To log in, make sure that the DOMAIN NAME is set to MATHSTAT.
- Use the workshop username: primesw The password will be distributed during the lab.
- On most of the machines, your work should be saved to the folder C:\temp (on some machines, this folder will be D:\temp).
- At the end of the lab, make sure that you logout from the computer (Start > Shutdown > Close all programs and logon as a different user.)
- These notes can be accessed in the folder primesw$ on `Powerdrift'[J]. Click on Start > Run and type

  j:\laboratory.pps and click on OK

To run the ITSM software:

- click on Start > Run and type
  j:\itsm2000\itsm.exe and click on OK
- Alternatively, click on the itsm icon in the
  primesw$ on `Powerdrift'[J] folder

DataFiles in ITSM

- Data file names should have the extension .TSM
- Data files should be located in the primesw$ on `Powerdrift'[J] folder.
- For univariate analyses, all data should be stored in a single column--each value must be on a separate line. For multivariate analyses, the m-variate time series must be stored in m columns.

Lab Tasks

1. This exercise relates to the maunaloa.tsm data consisting of monthly totals of CO_2 at Mauna Loa Oct `58 - Sept `90.

   • Open the maunaloa.tsm project.

   • Open the acf/pacf and the periodogram plot and the note their shape.

   • Select the option Transform > Classical and check the boxes for seasonal fit (period=12) and polynomial fit (quadratic trend). How do the plots of the acf/pacf and periodogram change after selecting these transformations?

   • Select Transform > Show Classical fit to assess the fit.

• Forecast 48 months ahead using the option Forecast > ARMA (Try setting the endpoint at 360.) Add 90% prediction bounds (enter WN variance as .272213) in the forecast dialog box.

• Try fitting the best ARMA(p,q) model (click on the blue AUT button, enter max AR order as 13 and max MA order as 5). Look at the acf/pacf of the residuals (the third green button). Do the residuals look uncorrelated. Test the residuals for randomness (Statistics > Residual Analysis > Tests of Randomness). What do these tell us (ask if you are not sure)?

• Are the mean square errors for prediction larger or smaller after fitting the ARMA model to the residuals?

• Simulate observations from the final model (Model > Simulate). Do the simulated observations look anything like the original data?

2. The NEE data (daily data from Jan 1, 1992 to Dec 31, 2001).

- Open the nee_24.tsm project.
- Open the acf/pacf and periodogram windows and note their appearance.
- Using the Transform > Classical option, remove seasonality (period 365) and a linear trend and note the changes in plots of the time series, acf/pacf and periodogram.
- Select Statistics > Residual Analysis > Tests of randomness option and comment on results.
- Select Transform > Show Classical fit to assess the fit.

- Try fitting the best ARMA(p,q) model (click on the blue AUT button, enter max AR order as 5 and max MA order as 5). Look at the acf/pacf of the residuals (the third green button). Do the residuals look uncorrelated. Test the residuals for randomness (Statistics > Residual Analysis > Tests of Randomness). What do these tell us (ask if you are not sure)?

- Simulate observations from the final model (Model > Simulate). Do the simulated observations look anything like the original data?

- Select Forecast > ARMA and forecast two years (730 days ahead) with endpoint set to 3288.

3. The NEE data (hourly data from Jan 1, 1993 to Dec 31, 1993).

- Open the nee_Jan-Dec93.tsm project.
- Open the acf/pacf and periodogram windows and comment on their appearance.
- Using the Transform > Classical option, remove seasonality (period 24) and a linear trend and note the changes in plots of the time series, acf/pacf and periodogram.
- Select Statistics > Residual Analysis > Tests of randomness option and comment on results.
- Select Transform > Show Classical fit to assess the fit.

- Try fitting the best AR(p) model (click on the blue PRE button, and check box Find AR model with min AICC). Look at the acf/pacf of the residuals (the third green button). Do the residuals look uncorrelated. Test the residuals for randomness (Statistics > Residual Analysis > Tests of Randomness). What do these tell us (ask if you are not sure)?

- Simulate observations from the final model (Model > Simulate).

Do the simulated observations look anything like the original data?

4. The tundra.tsm data file contains the average maximum temperature over the month of September for the years 1895-1993 in the areas of the US whose vegetation is classified as tundra.

- Open the tundra.tsm project.
- Open the acf/pacf comment on its appearance.
- Fit a straight line to the data by selecting the menu items Regression > Specify (check intercept term and enter 1 for order of polynomial), then click on the GLS button followed by the MLE button. Is the slope of the regression line significant?
- Try fitting the best ARMA(p,q) model (click on the blue AUT button, enter max AR order as 5 and max MA order as 5). Select Transform > Show Classical fit to assess the fit. Is the slope still significant?

5. Create a blank itsm project and generate 200 independent and identically distributed observations from a N(0,1) distribution (Model > simulate). Inspect the sample acf and histogram of the data.

6. Generate 200 independent and identically distributed observations from a Cauchy distribution. (At the Model > Specify step, click on the Change noise distribution button and select stable distribution with default parameters 1 and 0.) Compare the sample acf and histogram of this data set with that in 2.