Electricity demand

The series `vic_elec` is in the `tsibbledata` package (part of `fpp3`). It has electricity demand for Victoria, Australia, measured every half hour from 2012 to 2014.

1. Get started by converting to daily data. Use `index_by(Date)` and then summarize to get the total Demand for each day, the daily high temperature, and whether the day is a holiday.

2. Make a seasonal plot of Demand (with `gg_season`) to see the annual pattern of electricity use.

3. Since electricity use depends on heating, cooling, and the length of days, we concentrate only on the months June, July, August **(winter in Australia)**.

   Filter the data to just select those three months and just the year 2013. Make time plots of Demand and daily high temperature.

4. Make a plot showing Demand as a function of daily high temperature. What do you observe?

5. There are clearly two types of day in this data - days with and without work. Create a new variable that distinguishes the two types of day. Recreate your plot from (4) and color the points according to your new variable. Add regression lines for each group.

6. Create a time series linear model of Demand using high temperature and your workday variable. Check the coefficients. Describe how demand changes with a 1°C change in temperature. How much does a workday affect Demand?

7. Make a plot showing Demand as well as your fitted model. (You'll need to `augment()` your model.)

8. Plot the residuals of your model. What do you observe? Check if they are white noise using a Ljung-Box test.

Ocean temperatures

The series `sst.csv` contains weekly sea surface temperatures (°C) in four regions of the Southern Pacific ocean, from 1981 to 2023. These temperatures are indicators of the weather phenomenon known as El Niño.

1. Fit a linear model for the Nino4 region temperatures with time as the predictor. Plot the fitted values and data on the same chart. What does the trend coefficient tell you about ocean temperatures in this part of the ocean?

2. Fit a linear model and include trend and seasonal dummy variables. How many dummy variables did this require? Plot the fitted values and the data on the same chart.

3. Fit a linear model and include trend and Fourier series terms. Minimize cross validation error (CV) to select the best number of terms. Plot the fitted values and the data on the same chart.

4. Compare your two models (seasonal dummy and Fourier) visually and by using measures of predictive accuracy.