Quantitative Methods in Geography

Geo 340

William D. McCoy

Department of Geosciences
University of Massachusetts, Amherst

Autumn, 2009

Objectives of time series analysis

- Description of data
- Construction of a model to explain data
- Understand enough to forecast
- Understand enough to control the process

Time series components

- trend
- cyclic (period generally $>$ one year)
- seasonal (period equal to or less than one year)
- irregular or random

Convolution filters

A moving average applied to a time series is just one type of convolution filter. We can write the general form of a linear convolution filter (two-sided or non-causal) as

$$y_t = \sum_{j=-q/2}^{q/2} a_j x_{t-j},$$

where $t = 0, 1, 2, \ldots,$ q is even, and x_t is the unfiltered time series. The a_j are just a series of filter coefficients that are applied to the series. In the special case of a simple, unweighted moving average,

$$a_j = \frac{1}{q+1},$$

where $q + 1$ is equal to the length of the vector of filters.

Mauna Loa carbon dioxide example

To apply a convolution filter to a time series in **R**, use the function `filter()`. For example, get the `co2` dataset in **R** and plot it.

```
> data(co2)
> plot(co2)
```

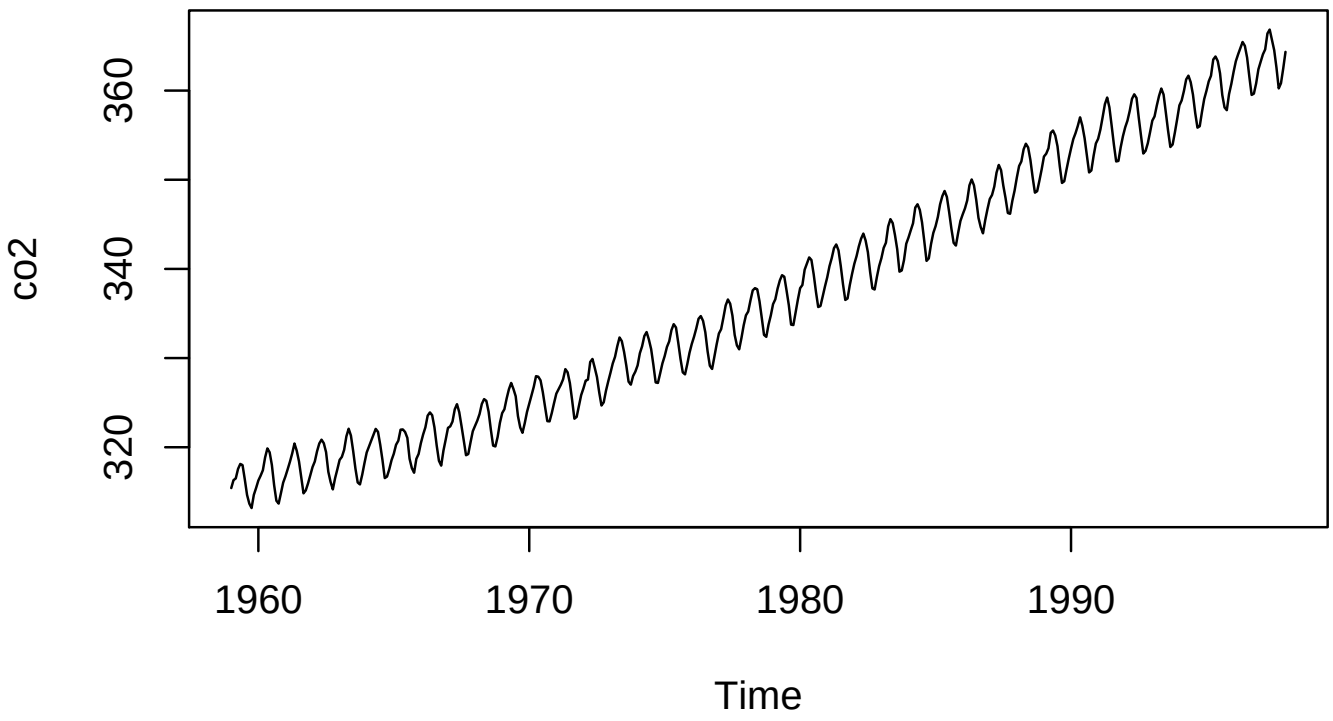

Applying a filter

Now apply a simple, unweighted, 12-month, moving average filter to the co2 data.

```
> co2.filtered <- filter(co2, c(1/24, rep(1/12, 11), 1/24))
> plot(co2)
> lines(co2.filtered, col = "red")
```

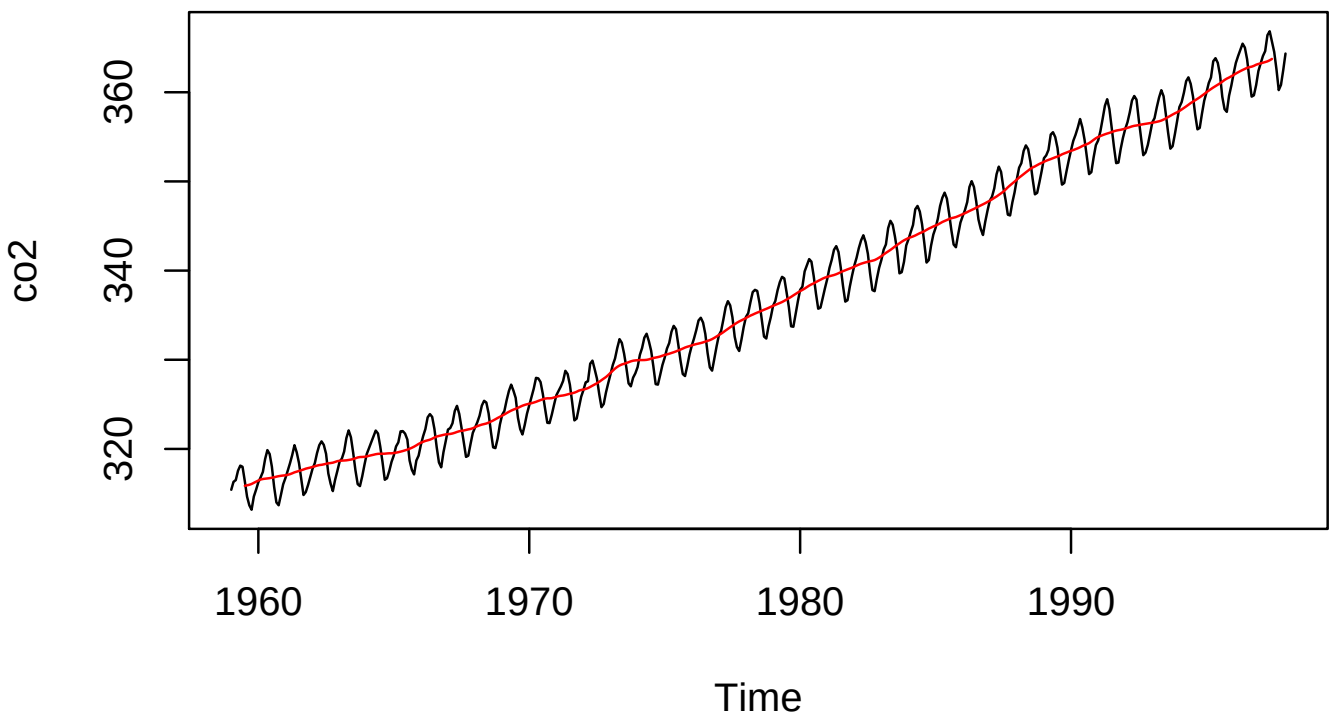

Removing the trend

```
> co2.seasonal <- co2 - co2.filtered
> plot(co2.seasonal)
```

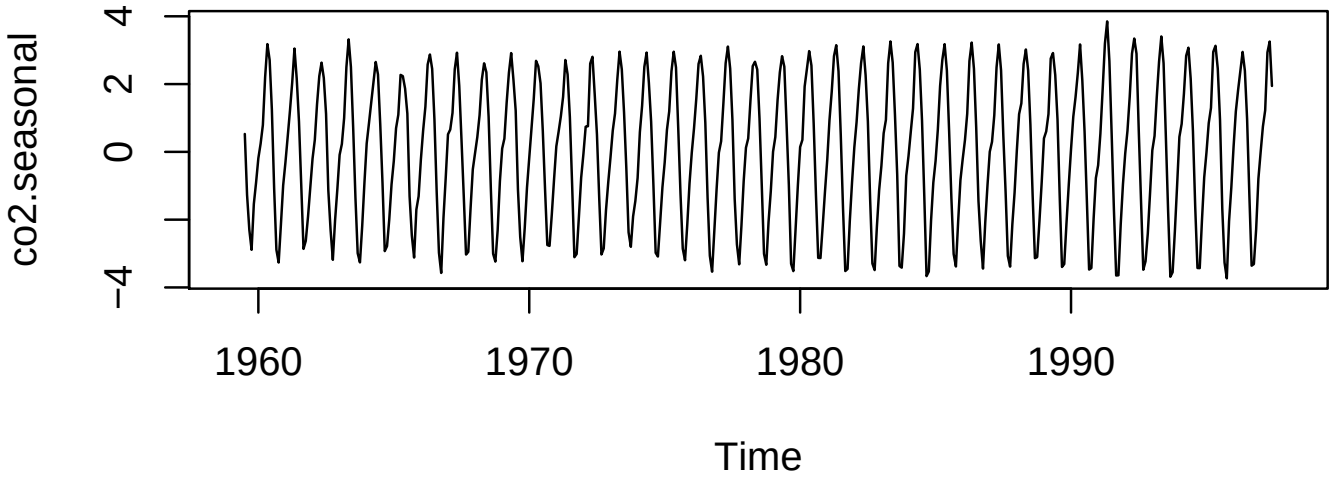

Decomposition of a time series

The **R** function `stl()` can be used to decompose a time series. The function first uses a locally weighted regression, `loess()`, to smooth each seasonal sub-series (*i.e.*, all the Januarys, all the Februarys, *etc.*). These are then removed from the series to reveal the trend. The overall mean is subtracted from the seasonal component and added to the trend. This whole process is repeated several times to smooth the components. Finally, both the trend and the seasonal components are subtracted from the original series to find the remainder. The remainder may contain both cyclic and random components.

```
> co2.stl <- stl(co2, s.window = 21)
> plot(co2.stl)
```

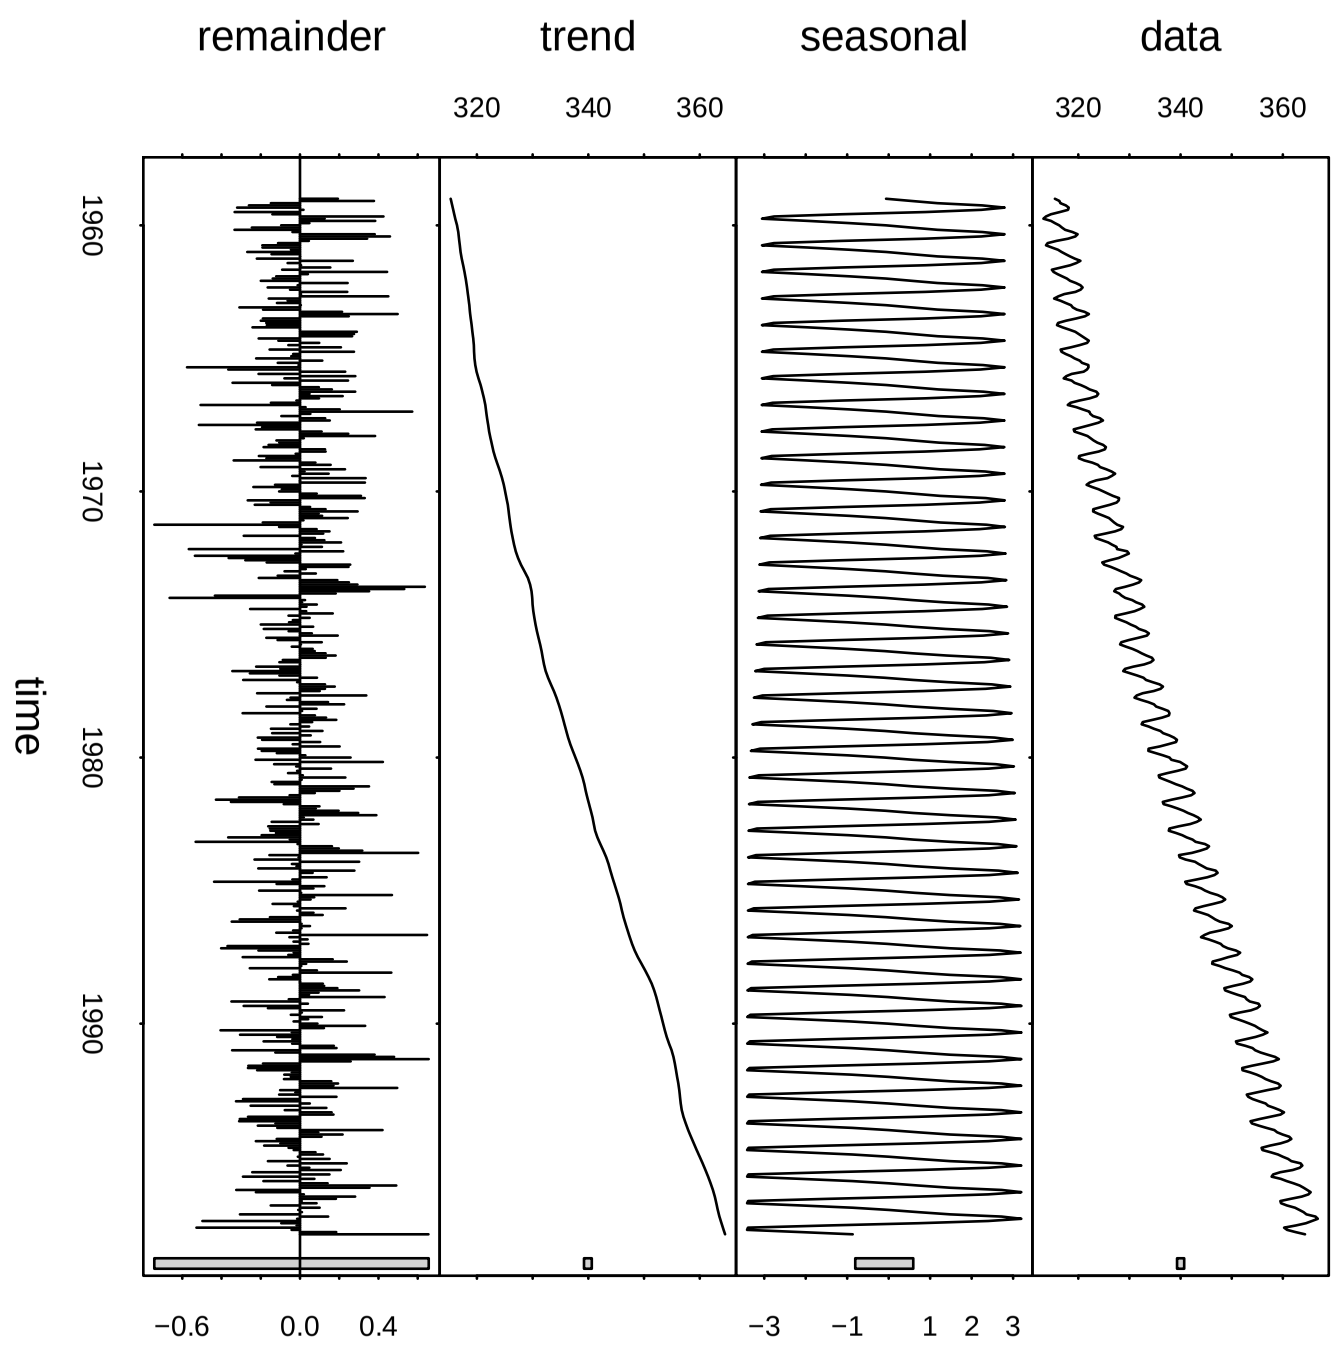
remainder
trend
seasonal
data
320
340
360
320
340
360
1960
1970
1980
1990
time
−0.6
0.0
0.4
−3
−1
1
2
3

```
> monthplot(co2.stl$time.series[, 1])
```