

Stat 565


Trend And Decompositions

Jan 14 2014

Charlotte Wickham

stat565.cwick.co.nz

Your turn



How could we get a feel for the long term trend in this series?

Exploring the trend

It's often hard to see a trend in the presence of noise or seasonality, some options:

Aggregate

use annual data instead

Smooth

moving average or add a line or curve

Subtract

remove noise/seasonality

Aggregate

To examine the long term trend we could average out the seasonality by averaging over the length of the seasonal cycle.

E.g. with temperature, average over a year.

We saw how to do this with plyr, on Thursday.

We could also achieve this same goal using a model.

Aggregate

We could fit a regression model with a fixed effect for each year, and examine those fixed effects,

$$\text{temp}_t = \mu_{\text{year}(t)} + \text{noise}$$

different mean
for each year

```
year_fit <- lm(temp ~ factor(year), data = corv,  
               na.action = na.exclude)
```


I want to treat year
like a category
not a number

Missing values are given
NA in fitted and residuals


Aggregate

```
corv$annual_avg <- fitted(year_fit)
```

```
qplot(year, annual_avg, data = corv,  
      geom = "line")
```



```
qplot(date, temp, data = corv,  
      geom = "line") +  
  geom_line(aes(y = annual_avg),  
            colour = "blue")
```



Smooth

Moving average `filter`

Linear regression `lm`

Local regression `loess`

many more ... (`gam`)

Smooth: Moving average *filter*

A moving average takes the average of n consecutive values.

If it is centered, take the average of $n+1$ consecutive values giving a half weight to numbers on each end. usually for even n

For daily data, $n = 7$, a weekly moving average

$n = 30$, approximately a monthly moving average

$n = 365$, approximately a yearly moving average

Smooth: Moving average *filter*

filter works on `ts` objects, so let's get our temperatures in to `ts` properly.

There were dates missing from our series:

```
all_days <- data.frame(date = seq(ymd("2000/01/01"),  
                                ymd("2011/12/31"), by = "day"))
```

```
# merge them together
```

```
corv <- join(corv, all_days)
```

```
# force it to be in time order
```

```
corv <- corv[order(corv$date), ]
```

```
# make a ts object
```


```
corv_ts <- ts(corv$temp, start = 2000, freq = 365.25)
```

Smooth: Moving average **filter**


```
n <- 7 # we have daily data so this is a week
corv$weekly_ma <- filter(corv_ts, filter = rep(1, n)/n)
n <- 30 # approximately a month
corv$monthly_ma <- filter(corv_ts, filter = c(1/2, rep(1, n-1), 1/2)/n)
n <- 365 # approximately a year
corv$annual_ma <- filter(corv_ts, filter = rep(1, n)/n)

qplot(date, temp, data = corv, geom = "line", alpha = I(0.5)) +
  geom_line(aes(y = weekly_ma, colour = "weekly_ma"), size = 1) +
  geom_line(aes(y = monthly_ma, colour = "monthly_ma"), size = 1) +
  geom_line(aes(y = annual_ma, colour = "annual_ma"), size = 1) +
  scale_colour_brewer("Moving average", pal = "Set1") +
  big_font
```

Smooth: Moving average *filter*




Smooth: Linear regression `lm`




```
qplot(date, temp, data = corv, geom = "line") +  
  geom_smooth(method = "lm", se = FALSE)
```

Smooth: Local regression **loess**



```
qplot(date, temp, data = corv, geom = "line") +  
  geom_smooth(method = "loess")
```


Smooth: Local regression **loess**



```
qplot(date, temp, data = corv, geom = "line") +  
  geom_smooth(method = "loess", span = 0.2)
```

Subtract

```
month_fit <- lm(temp ~ factor(month), data = corv, na.action = na.exclude)
corv$month_avg <- predict(month_fit)
corv$res <- residuals(month_fit)
```



```
qplot(date, month_avg, data = corv, geom = "line", group = year)
qplot(date, res, data = corv, geom = "line") +
  geom_smooth(method = "loess", span = 0.2)
```

A common decomposition

$$X_t = m_t + S_t + Z_t$$

Variable
measured
at time t

Trend

Seasonality

Noise

$$X_t = m_t S_t Z_t$$

$$X_t = m_t S_t + Z_t$$

some multiplicative analogs

Suggests a general approach

Describe and model the most obvious part

Subtract it from the series and repeat

When there is just "noise" left, examine it's variance and correlation

Hints for subtracting

Find a model for the part you want to subtract using: `lm`, `loess`, ...

Fit the model (with `na.action = na.exclude`)


Use `predict` to get the model prediction and `residuals` to do the subtraction

If you want to use a moving average, the output from `filter` is the prediction, do the subtraction "by hand"

An example


Corvallis data,
the largest
variation is
from the
seasonal
pattern.

Might be nice
to use a
smooth.




```
qplot(yday, temp, data = corv, geom = "line", group = year, alpha = I(.3)) +  
  geom_smooth(method = "loess", aes(group = 1), size = 1)
```

```
qplot(yday, seas_fit, data = corv,  
      geom = "line")
```



```
seas_fit <- loess(temp ~ yday, data = corv,  
                 na.action = na.exclude)  
corv$seas_fit <- predict(seas_fit,  
                        newdata = corv)  
corv$deseas <- residuals(seas_fit)
```

```
qplot(date, deseas, data = corv, geom = "line") +  
  geom_smooth(method = "loess", span = 0.2)
```



Your turn




Where would you start with this series?

Two automatic approaches

`stl` iteration of local regression
smooths


`decompose` iteration of moving averages
neither like missing values

```
plot(stl(corv_ts2, 365.25))
```



```
plot(decompose(corv_ts2))
```

Decomposition of additive time series



Next time...

Examining variance and correlation