

An introduction to GARMA models

Introduction to long memory models and GARMA models.

GARMA models are a type of time series models with special properties. These models are related to the “arima” models which can be fit in R by the standard “arima” function or the “Arima” function in the **forecast** package.

IF you are not very familiar with ARIMA models, it might be advisable to first have a look at the excellent online book by Rob Hyndman and George Athanasopoulos:- Forecasting: Principals and Practice.

GARMA models are a type of **long memory** model, also known as **fractal** models or **fractionally differenced** models. The reason for this is that they can model proceses which have a high correlation between observations which are far apart in time.

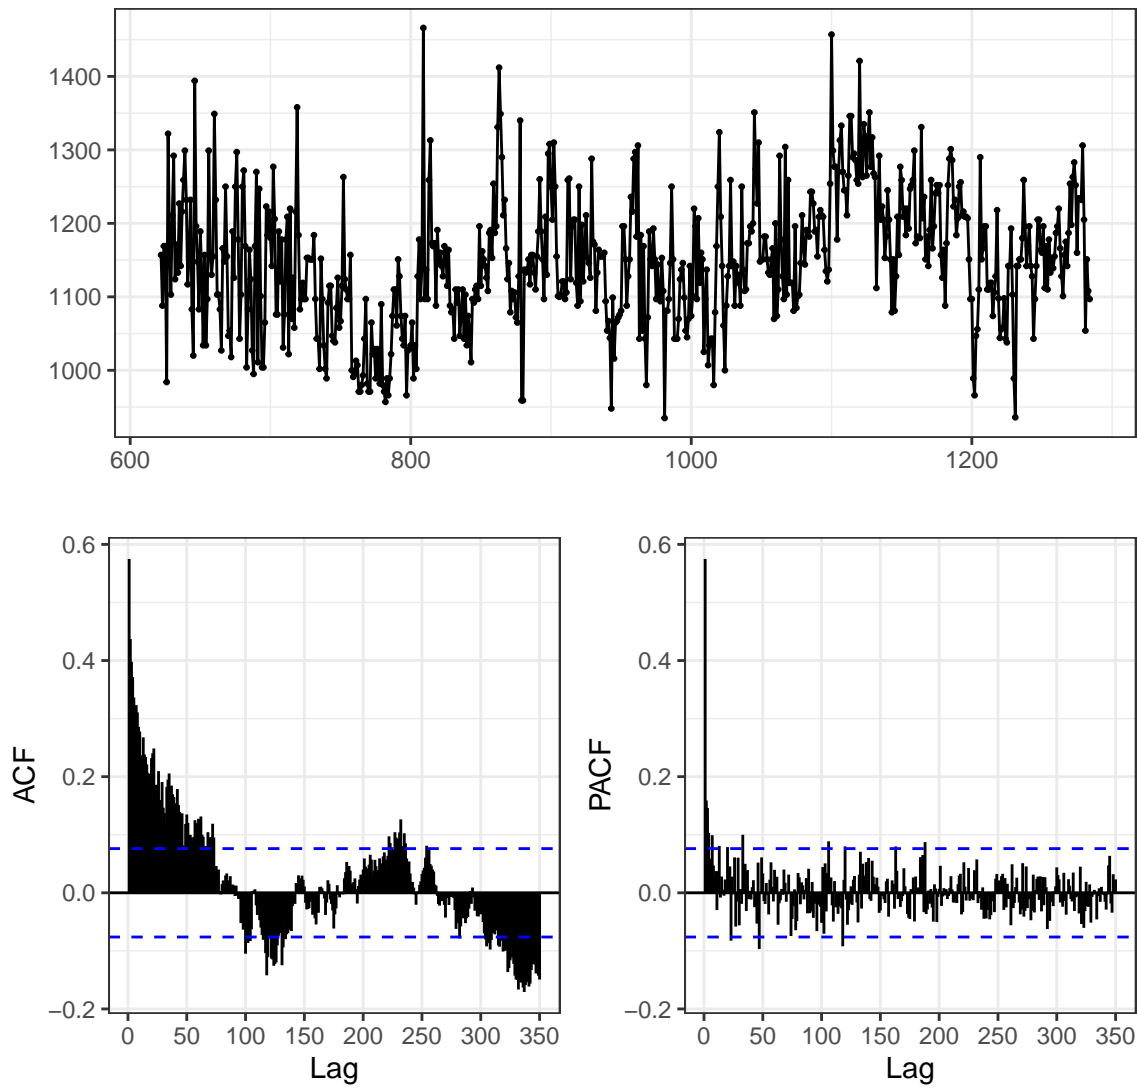
For instance, consider the data on the minimum level of the Nile (Tousson 1925), as measured from 622 to 1281 AD (from the **longmemo** package):

```
library(tidyverse)
library(forecast)

data(NileMin,package='longmemo')
# we'll just set the correct start year, for display purposes.
NileMin<-ts(as.numeric(NileMin),start=622,frequency=1)

ggtsdisplay(NileMin,lag.max=350, main='NileMin time series.', theme=theme_bw())
```

NileMin time series.



The first thing to note here is that the ACF (autocorrelation function) is showing highly significant correlations between Nile measurements, even at 350 years apart!

We won't attempt to address the question of how it is that a flood or minima level of a river flow today could be significantly related to what was happening 350 years ago, but scientists believe this is a real affect, not an artifact of measurement. This time series has been investigated many times by researchers.

The phenomena of long memory has been identified in a number of different types of time series from sunspots (Gray, Zhang, and Woodward 1989), the timing of xrays from Cygnus X1 (Greenhough et al. 2002), inflation rates, exchange rates (Caporale and Gil-Alana 2014) and even in non-coding DNA (Lopes and Nunes 2006).

ARMA models which can model long memory like this are sometimes known generically as 'long memory models' or 'fractionally differenced' models, and the default version of these can be fit by a number of existing R packages, including **fracdiff**, **longmemo**, **rugarch** and even **forecast**, among many others.

You might be wondering why these models might be useful - the answer is that they tend to produce long range forecasts which are more accurate than other short-memory models.

At time of writing there are few packages on CRAN which can fit GARMA models.

Package **tswge** provides a function “`est.garma.wge`” which can fit a GARMA model, but unfortunately can only do so via a grid search technique (searching a grid of values of the likelihood function), which tends to be slow and can result in rather inaccurate results. It also does not allow any MA terms to be fitted.

Package **waveslim** has functions “`spp.mle`” and “`spp2.mle`” which can identify the parameters of a white noise process only, using a wavelet method.


Package **VGAM** to have a function called “`garma`” but this actually fits very different models - it fits short-memory models which can include count data. Package **gsarima** also fits these short memory count models.

The models fit by the **garma** package are models which are not just long memory models, but where the long memory phenomena actually cycles in strength. An example of this is the Southern Oscillation Index.

```
soi <-
  read_fwf('soi.long.data',fwf_widths(c(5, rep(7,12)), c("year", 1:12)),col_types='nnnnnnnnnnnnnn') %>%
  gather(mth_no,soi,-year) %>%
  mutate(mth = lubridate::make_date(year,as.numeric(mth_no),1)) %>%
  select(mth,soi) %>%
  arrange(mth)

soi_ts <- ts(soi$soi,start=c(1951,2),frequency=12)

ggtsdisplay(soi_ts, lag.max=400, main='Southern Oscillation Index', theme=theme_bw())
```



In this plot the ACF shows the cyclical nature which is so typical of the GARMA models.

GARMA stands for Gegenbauer AR-MA, and the AR and MA letters have their usual meaning in time series of “auto-regressive” and “moving average”. ‘Gegenbauer’ refers to a mathematician Leopold Gegenbauer who lived in the 19th century and developed a series of orthogonal polynomials now called Gegenbauer


polynomials, and sometimes called ‘ultraspherical’ polynomials (Szegò 1959). These polynomials feature strongly in the mathematics of GARMA models (Gray, Zhang, and Woodward 1989), (Gray, Zhang, and Woodward 1994).

All long memory models, including GARMA models, share one thing in common. When you examine them on the frequency domain, they will show evidence of an unbounded (read: very large) peak in the spectrum. The more traditional models - known as “short memory” models do not have such a peak - whilst the spectrum may vary up and down, it has absolute bounds beyond which it will not go.

```
spectrum_nilemin <- spectrum(NileMin, plot=FALSE)
spectrum_soil    <- spectrum(soi_ts, plot=FALSE)

# now munge these lists together into a single dataframe.
spec_df <- rbind(data.frame(freq=spectrum_nilemin$freq,
                             spec=spectrum_nilemin$spec,
                             process='NileMin'),
                 data.frame(freq=spectrum_soil$freq,
                             spec=spectrum_soil$spec,
                             process='SOI'))

# and plot
ggplot(spec_df, aes(x=freq,y=spec)) +
  geom_line() +
  facet_wrap(~process,scales='free_y') +
  ggtitle('Spectrum of NileMin and SOI') +
  ylab('Intensity') +
  xlab(bquote('Frequency (0 -' ~ pi ~')')) + xlim(0,pi) +
  theme_bw()
#> Warning: Removed 447 row(s) containing missing values (geom_path).
```



From the above you can see that the NileMin spectrum shows an (essentially) unbounded peak at 0 which is the marker of a traditional long memory process.

However the SOI has 3 main peaks, all separated away from 0. Although these are not as large as the NileMin peaks, compared with the rest of the spectrum they are quite large, so it is not inappropriate to at least try to model the SOI by a 3-factor (or $k=3$) Gegenbauer model. This was first pointed out by (Lustig, Charlot, and Marimoutou 2017).

Technical details of the model.

The GARMA model as fit by the **garma** package is specified as

$$\phi(B) \prod_{i=1}^k (1 - 2u_i B + B^2)^{d_i} (X_t - \mu) = \theta(B) \epsilon_t$$

where

1. $\phi(B)$ represents the short-memory Autoregressive component of order p ,
2. $\theta(B)$ represents the short-memory Moving Average component of order q ,
3. $(1 - 2u_i B + B^2)^{d_i}$ represents the long-memory Gegenbauer component (there may in general be k of these),
4. X_t represents the observed process,
5. ϵ_t represents the random component of the model - these are assumed to be uncorrelated but identically distributed variates. Generally the routines in this package will work best if these have an approximate Gaussian distribution.
6. B represents the Backshift operator, defined by $BX_t = X_{t-1}$.

When $k=0$, then this is just a short memory model as fit by the stats “arima” function.

Fitting a short memory model

We have deliberately kept the fitting process close to that of the “arima” and “forecast::Arima” functions.

To illustrate basic usage of the routine, we will first look at fitting a simple ARIMA model to the “AirPassengers” data supplied with R. To achieve stationarity with this data, we’ll need to seasonally difference it. the “arima” function can do this but unfortunately “garma” as yet does not fit a seasonal model - seasonality is essentially modelled by the Gegenbauer components, but we won’t initially use that.

```
library(garma)
data(AirPassengers)

ap <- as.numeric(diff(AirPassengers,12))

# Arima model
arima_mdl <- arima(ap,order=c(9,1,0))
summary(arima_mdl)
#>
#> Call:
#> arima(x = ap, order = c(9, 1, 0))
#>
#> Coefficients:
#>          ar1      ar2      ar3      ar4      ar5      ar6      ar7      ar8
#>      -0.3157  -0.0131  -0.1354  -0.2238  -0.0083   0.0278  -0.1435  -0.0859
#> s.e.    0.0856   0.0915   0.0903   0.0914   0.0947   0.0939   0.0918   0.0924
#>          ar9
#>      0.2049
#> s.e.    0.0940
```

```

#>
#> sigma^2 estimated as 123.3: log likelihood = -501.71, aic = 1023.42
#>
#> Training set error measures:
#>
#>      ME      RMSE      MAE  MPE MAPE      MASE      ACF1
#> Training set 0.4280725 11.06271 8.690653 -Inf Inf 0.9225896 0.01427067

# GARMA model
# Note in the below we specify k=0.
# This tells the routine is not to fit a Gegenbauer/GARMA model.
garma_md1 <- garma(ap,order=c(9,1,0),k=0)
summary(garma_md1)
#>
#> Call:
#> garma(x = ap, order = c(9, 1, 0), k = 0)
#>
#> Coefficients:
#>      ar1      ar2      ar3      ar4      ar5      ar6      ar7      ar8
#>    -0.3169 -0.0309 -0.1393 -0.2034 -0.0111 -0.0221 -0.1289 -0.0895
#> s.e.    0.0954  0.0998  0.0991  0.0999  0.1018  0.0999  0.0990  0.0998
#>      ar9
#>    0.1480
#> s.e.    0.0954
#>
#>
#>
#> sigma^2 estimated as 125.3302: log likelihood = 147.443376, aic = -274.886752

```

As can be seen above, the coefficients produced are similar but not identical - the log-likelihood from the “garma” run is considerably lower than that produced by the “arima” run, indicating that the routine has in fact found a better solution (estimating parameters like these involves non-linear optimisation - “arima” uses R’s built-in optimiser called “optim”; garma however by default uses “solnp” from the **Rsolnp** package).

Fitting a GARMA model.

In this section we’ll look at fitting a GARMA model to the Sunspots data, again as supplied with R. This data has been analysed many times in the literature; the first time was by (Gray, Zhang, and Woodward 1989). Generally, authors have used a standard subset of this data from 1749 to 1924.

The Sunspot data consists of counts of sunspots as observed over a considerable period. In recent years a large project by the Royal Observatory of Belgium has checked and extended this key dataset (Clette et al. 2014), but we’ll use the original here.

```

library(garma)

data(sunspot.year)

# Next we subset the data to ensure we are using the years used in the literature.
sunspots <- ts(sunspot.year[49:224],start=1749,end=1924)

# Now as in Gray et al 1989 we fit a GARMA(1,0) model:
sunspots_garma_md1 <- garma(sunspots, order=c(1,0,0),k=1,method='CSS')

summary(sunspots_garma_md1)

```

```

#>
#> Call:
#> garma(x = sunspots, order = c(1, 0, 0), k = 1, method = "CSS")
#>
#> Coefficients:
#>      intercept      u1      fd1      ar1
#>      45.1371  0.8475  0.4228  0.4937
#> s.e.      3.8143  0.0070  0.0649  0.1042
#>
#>
#>      Factor1
#> Gegenbauer frequency:  0.0891
#> Gegenbauer Period:    11.2292
#> Gegenbauer Exponent:   0.4228
#>
#>
#> sigma^2 estimated as 228.0985: part log likelihood = -639.565233

```

Above, we have specified **method**='CSS' to ensure we are using a method as close as possible to that used by (Gray, Zhang, and Woodward 1989). The “garma” function uses a frequency-domain method known as the “whittle” method by default, since this method not only produces very accurate results very quickly, but has a lot of theoretical results available to support its use - for example (Giraitis, Hidalgo, and Robinson 2001).

The following table compares the values found by Gray et al and by the “garma” function:

Parameter	Gray et al	garma
intercept	44.78	45.1371
u	0.85	0.8475
d	0.42	0.4228
ar1	0.49	0.4937

As you can see, the results are quite close (Gray et al only published their results to 2 decimal places).

Notice also that the routine displays the Gegenbauer Period - in this case 11.2 years - which corresponds nicely with the “known” 11 year sunspot cycle.

Also shown is the degree of fractional differencing and the Fractional Dimension of the original series.

Whilst this model does fit the data quite well, is quite simple and does not produce very effective forecasts. A better model is the GARMA(8,0) model, which was also examined by (Gray, Zhang, and Woodward 1989) and which we fit then forecast below, for the 11 years of a sunspot cycle.

```

# fit a GARMA(8,0,0) model
sunspots_garma_mdl <- garma(sunspots, order=c(8,0,0),k=1,method='CSS')
summary(sunspots_garma_mdl)
#>
#> Call:
#> garma(x = sunspots, order = c(8, 0, 0), k = 1, method = "CSS")
#>
#> Coefficients:
#>      intercept      u1      fd1      ar1      ar2      ar3      ar4      ar5
#>      44.6954  0.8511  0.3337  0.6599 -0.1158 -0.1159  0.1218 -0.0489
#> s.e.      3.3887  0.0052  0.1538  0.2786  0.1354  0.0896  0.0868  0.1052
#>
#>      ar6      ar7      ar8
#>      0.0811 -0.0838  0.2062
#> s.e.  0.0904  0.0936  0.0772

```

```


#>
#>                                     Factor1
#> Gegenbauer frequency:      0.0880
#> Gegenbauer Period:       11.3678
#> Gegenbauer Exponent:      0.3337
#>
#>
#> sigma^2 estimated as 211.0765: part log likelihood = -632.741046

#prepare 'future' actuals data for plotting
future_df <- data.frame(yr=1925:1935, sunspots=sunspot.year[225:235],grp='Future Actuals')

ggplot(sunspots_garma_mdl, h=11) +
  geom_line(data=future_df,aes(x=yr,y=sunspots)) +
  ggtitle('Sunspot Forecast using GARMA(8,0)')

```

Sunspot Forecast using GARMA(8,0)



References

- Caporale, G, and L Gil-Alana. 2014. "Long-Run and Cyclical Dynamics in the Us Stock Market." *Journal of Forecasting* 33: 147–61.
- Clette, F, L Svalgaard, J Vaquero, and E Cliver. 2014. "Revisiting the Sunspot Number - a 400 Year Perspective on the Solar Cycle." *Space Science Reviews*. 186 (1-4): 35–103.
- Giraitis, L, J Hidalgo, and P Robinson. 2001. "Gaussian Estimation of Parametric Spectral Density with

Unknown Pole.” *The Annals of Statistics* 29 (4): 987–1023.

Gray, H, N Zhang, and W Woodward. 1989. “On Generalized Fractional Processes.” *Journal of Time Series Analysis* 10 (3): 233–57.

———. 1994. “On Generalized Fractional Processes - a Correction.” *Journal of Time Series Analysis* 15 (5): 561–62.

Greenhough, J, S Chapman, S Chaty, R Dendy, and G Rowlands. 2002. “Characterising Anomalous Transport in Accretion Disks from X-Ray Observations.” *Astronomy and Astrophysics* 385: 693–700.

Lopes, S, and M Nunes. 2006. “Long Memory Analysis in Dna Sequences.” *Physica A* 361: 569–88.

Lustig, A, P Charlot, and V Marimoutou. 2017. “The Memory of Enso Revisited by a 2-Factor Gegenbauer Process.” *International Journal of Climatology* 37 (5): 2295–2303.

Szegò, G. 1959. *Orthogonal Polynomials*. AMS, New York.

Tousson, O. 1925. *Mémoire Sur L'Histoire Du Nil; Mémoire de L'Institut d'Egypte*.