An Efficient Method for Characterising Noise in the Time Domain

Jim Barrett, Will Farr
School of Physics and Astronomy, University of Birmingham

Coloured Noise

Noise is an unpredictable perturbation to a measurement of a physical process. It is usually impossible to avoid noise completely when designing an experiment, so it is important to understand its contribution.

In an ideal world, all noise would be white. White noise, by definition, is completely statistically independent at all points in time, so that noise contributions in the past have absolutely no bearing on the present. In this regime, one can often make a robust analysis of the underlying process.

However, in reality, most noise is coloured. Coloured noise has non-zero correlations in time, meaning that the noise on a measurement at the present time has some dependence on previous noise contributions. This poster presents an efficient method for characterising the noise in unevenly sampled, time domain data.

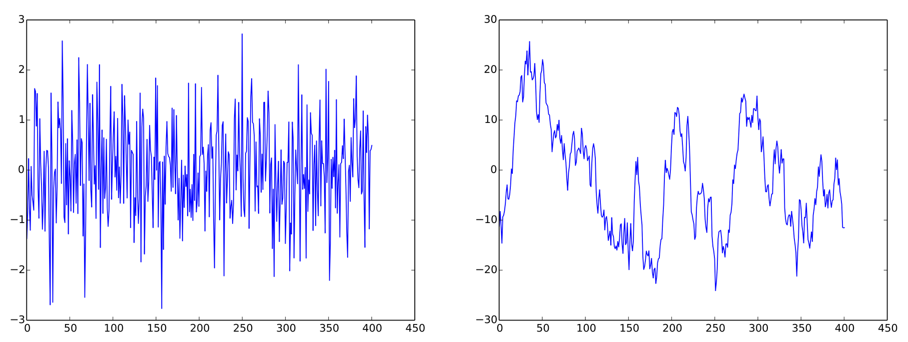

(a) White Noise (b) Coloured Noise

Figure: Realisations of noise with (a) no correlations and (b) correlations on timescales of 50 time units

Power Spectral Density

One of the most common and useful ways of visualising timescales and correlations in noise is by plotting the power spectral density (PSD). Intuitively, the PSD can be thought of as a type of Fourier analysis, giving the power in a stochastic process as a function of frequency.

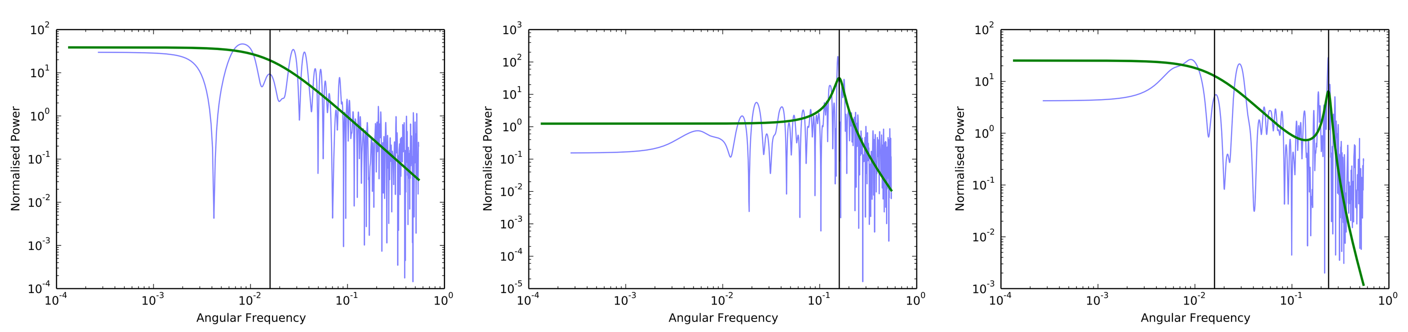

(a) 'Knee' feature (b) 'Peak' feature (c) Combined PSD

Figure: PSDs of 3 different noise realisations with different features. The faint blue gives the Lomb-Scargle periodogram, the green gives the 'true' PSD and the black lines guide the eye for the injected features.

The qualitative features of a PSD can be an indicator for the properties of the represented noise. Figure (a) gives the PSD of a noise process correlated over a decaying timescale. Figure (b) gives a PSD for noise with a periodic timescale and figure (c) shows a PSD with both of these features. Quantifying all these features for a PSD is equivalent to characterising the represented noise process itself.

CARMA Processes

A continuous, auto-regressive moving average (CARMA) process is one where at any given time the value of a measurement can be described by a linear combination of previous measurements and noise contributions. Algebraically, this can be generally expressed as a linear ordinary differential equation.

$$\prod_{i=1}^{p}\left[\frac{d}{dt}-r_i\right]y(t)=\prod_{j=1}^{q}\left[\frac{d}{d\xi}-b_i\right]\eta(\xi)$$

The $\{r_i\}$ have units of inverse time, and can be interpreted as the timescales over which the observed noise in the system is correlated. The $\{b_i\}$ also have units of inverse time and characterise the timescales over which the unobserved white noise error terms $\eta(\xi)$ contribute to the measurement.

A key property of auto-regression is that any given 'measurement' of a noisy process $y(t_k)$ can be written as a linear sum of some number of previous measurements and a separate sum of purely stochastic terms.

$$y(t_k)=\underbrace{\sum_{i=1}^{k-1}\alpha_i y(t_i)}_{\text{deterministic}}+\underbrace{\sum_{i=1}^{k-1}\beta_i}_{\text{stochastic}}$$

One can naively calculate a dense covariance matrix by considering $\mathbf{C_{ij}}=\langle y(t_i)y(t_j)\rangle$. However, if we are able to employ the transformation

$$x(t_k)=y(t_k)-\sum_{i=1}^{k-1}\alpha_i y(t_i)=\sum_{i=1}^{k-1}\beta_i$$

Then the covariance matrix in the x basis is sparse, and the expense of inverting it for the purposes of likelihood computation can be improved from $O(N^3)\rightarrow O(N)$.

Relating CARMA Models to PSDs

The likelihood described above can be naturally parametrised in terms of the variance of the driving noise $\sigma^2=\mathbf{Var}(\eta(\xi))$ and the set of auto-regressive (AR) and moving average (MA) roots. The resultant PSD can be written in terms of these parameters.

$$\mathbf{PSD}=\sigma^2\frac{\left|\prod_{j=1}^{q}(2\pi i f-b_j)\right|^2}{\left|\prod_{j=1}^{p}(2\pi i f-r_j)\right|^2}$$

The quantitative features of the PSD can be predicted from the values of the AR roots. Where two AR roots appear in a complex conjugate pair, they form a 'peak' feature in the PSD with centroid $\frac{|\mathbf{Im}(r)|}{2\pi}$. A real AR root corresponds to a 'knee' feature at $\frac{1}{|\mathbf{Re}(r)|}$.

Applying CARMA Models - XB158

A binary system is a pair of massive objects orbiting their common centre of mass. X-Ray binaries are a particular class of binary system which is luminous in X-Rays, and are thought to typically consist of a compact object, such as a neutron star or black hole, and a more normal star. X-Rays are emitted as matter is dragged from the normal star on to the compact companion.

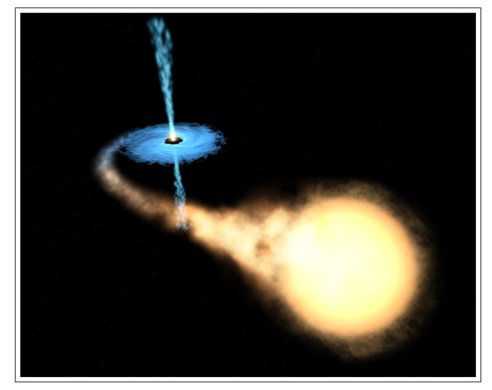

Figure: An artist's impression of an X-Ray binary, with a compact object accreting matter from its companion

Recent observations of one particular X-Ray binary (XB158) highlighted the possibility of a 'super-orbital' period on a timescale of approximately 5.65 ± 0.05 days (arXiv:1501.01978). We used Bayesian inference to fit a CARMA model to the data.

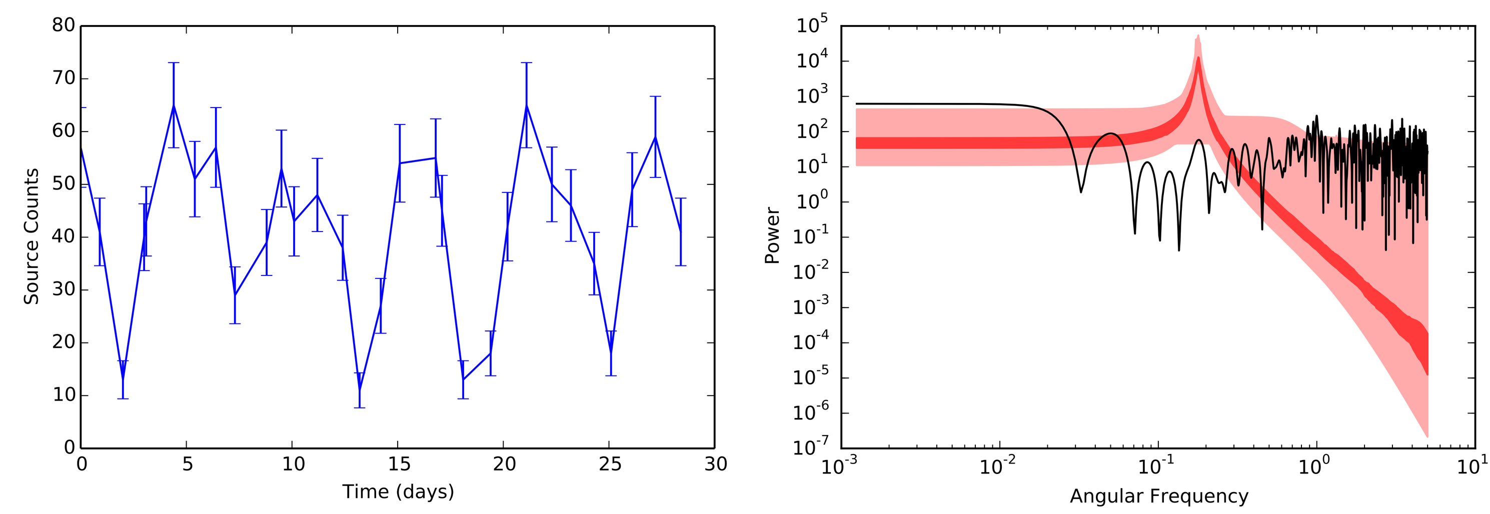

(a) Observations (b) Power Spectral Density

Figure: (a) The raw observations (b) The power spectral density. In black is the Lomb-Scargle periodogram to guide the eye, then in red are the 1σ and 3σ percentiles for the CARMA model.

Figure (b) shows a very clear 'peak' feature in the PSD, indicative of a periodic signal embedded in the data. Further investigation of the relevant parameters reveals a posterior on the timescale of this periodic behaviour in exact agreement with the original published result.

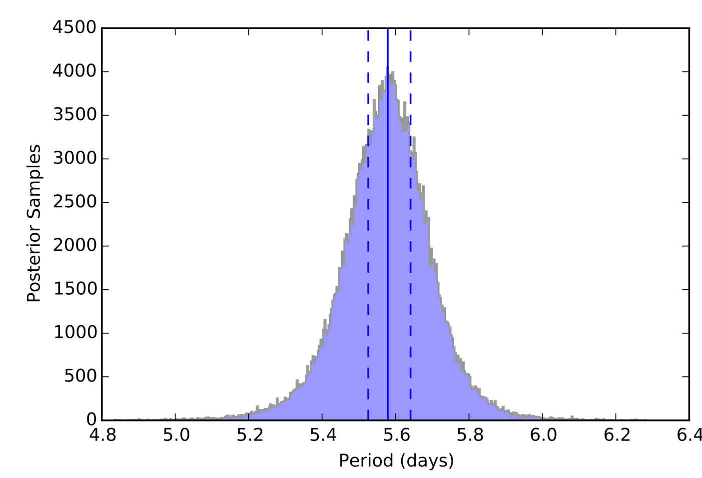

Figure: A Histogram of $\sim 90\%$ of the posterior probability for the period of a signal found in the XB158 dataset. The solid blue line gives the median and the dashed lines the 1σ confidence interval.