An Efficient Method for Characterising Noise in the Time Domain

Jim Barrett, University of Birmingham

6th February 2015

Stochastic Processes

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 'whitening' noise in the presence of unevenly sampled data.

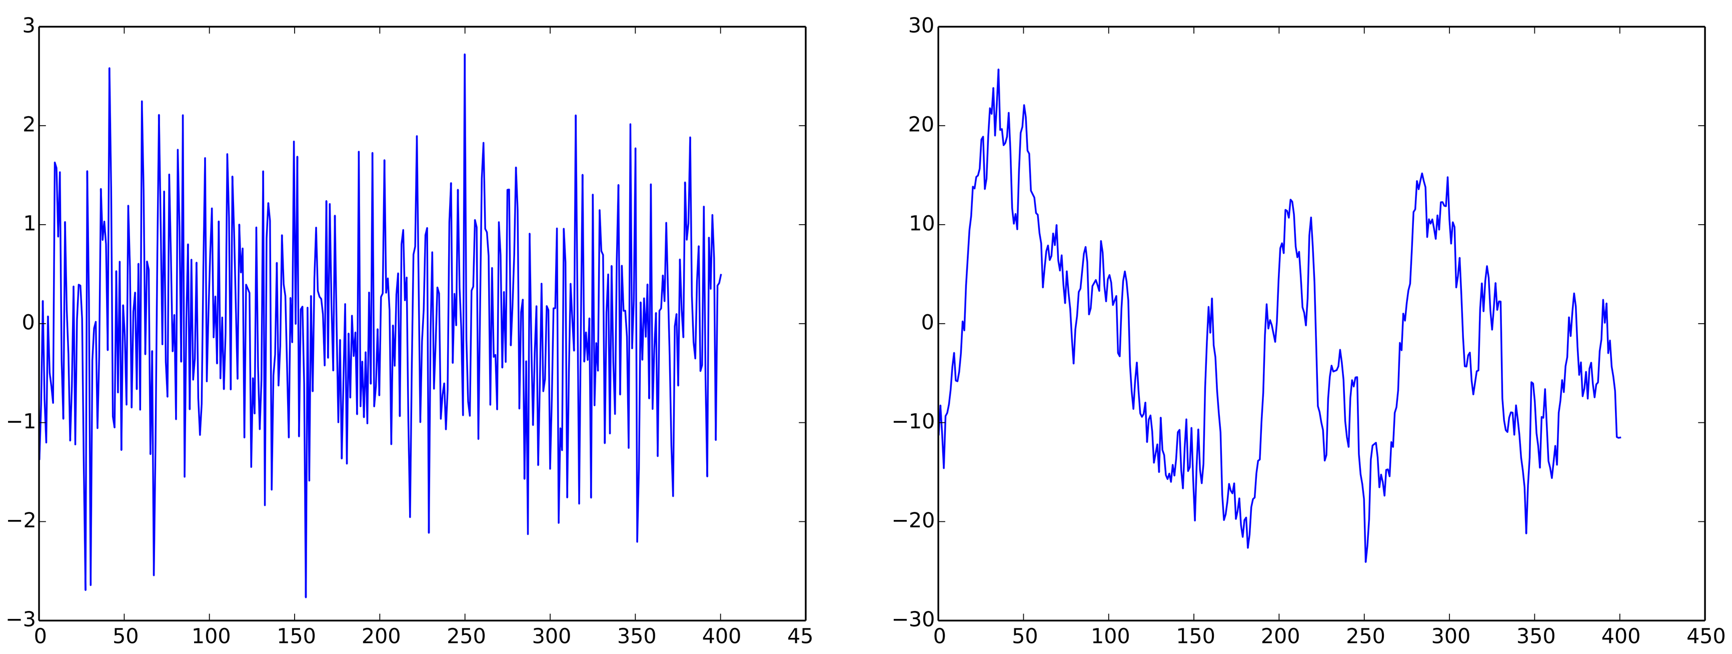

(a) White Noise (b) Coloured Noise

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

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.

In general, given a noisy time series $[y(t_0), y(t_1), y(t_2)\ldots]$ and assuming that the noise is normally distributed in every parameter, it is possible to write down the relative likelihood function $(\mathcal{L})$

$$\log(\mathcal{L}) = -\frac{1}{2}|C| - \frac{1}{2}y^T C^{-1} y$$

Where C is the covariance matrix for the model, $C_{ij} = \langle y(t_i)y(t_j)\rangle$, depending on the set of parameters $\{r, b\}$. The covariance is a square matrix with the same length as the number of measurements, and inverting it is an expensive process $(O(n^3))$.

By exploiting the properties of CARMA models, it is possible to isolate the deterministic (physically driven) and stochastic (noise driven) contributions to any given measured value.

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

So that by employing the transformation

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

The covariance matrix in the x basis takes on a banded form, and the resulting likelihood can be found with relatively little computational expense $(O(n))$.

Bayesian Analysis

Baye's theorem relates the probability density that the parameters governing a model, given the observed data and the form of the model $P(\theta|D,M)$ (Posterior), to the product of the probability of obtaining the data, given the parameters and the model $P(D|M,\theta)$ (Likelihood), and any other prior information $P(\theta|M)$ (Prior).

$$P(\theta|D,M) \propto P(D|M,\theta)\cdot P(\theta|M)$$

By exploring the space of the parameters θ that characterise the model, it is possible to find the values of the parameters for which the model is most likely to correctly predict the observations. This is known as 'sampling' the parameter space.

An Application - Transits

A transit signal is the dimming of the light received from a star when one of its orbiting planets passes between it and the observer. The CARMA method described in this poster can be used to whiten the correlations in a noisy transit signal.

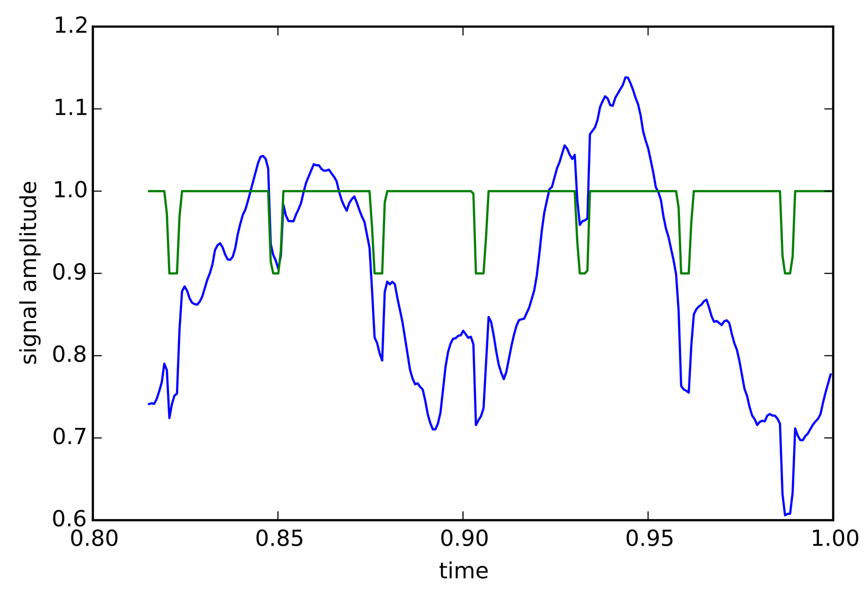

Figure: A perfect transit signal (green) with the same signal added to some noise with correlation timescales comparable to the transit parameters (blue). The amplitude of the noise is such that the signal to noise ratio is ~ 90.

A CARMA representation of the transit is sampled over the characteristic parameters $\{r, b\}$ until a distribution of the most likely combinations is found. The most likely values can then be used to filter out correlations, leaving IID white noise.

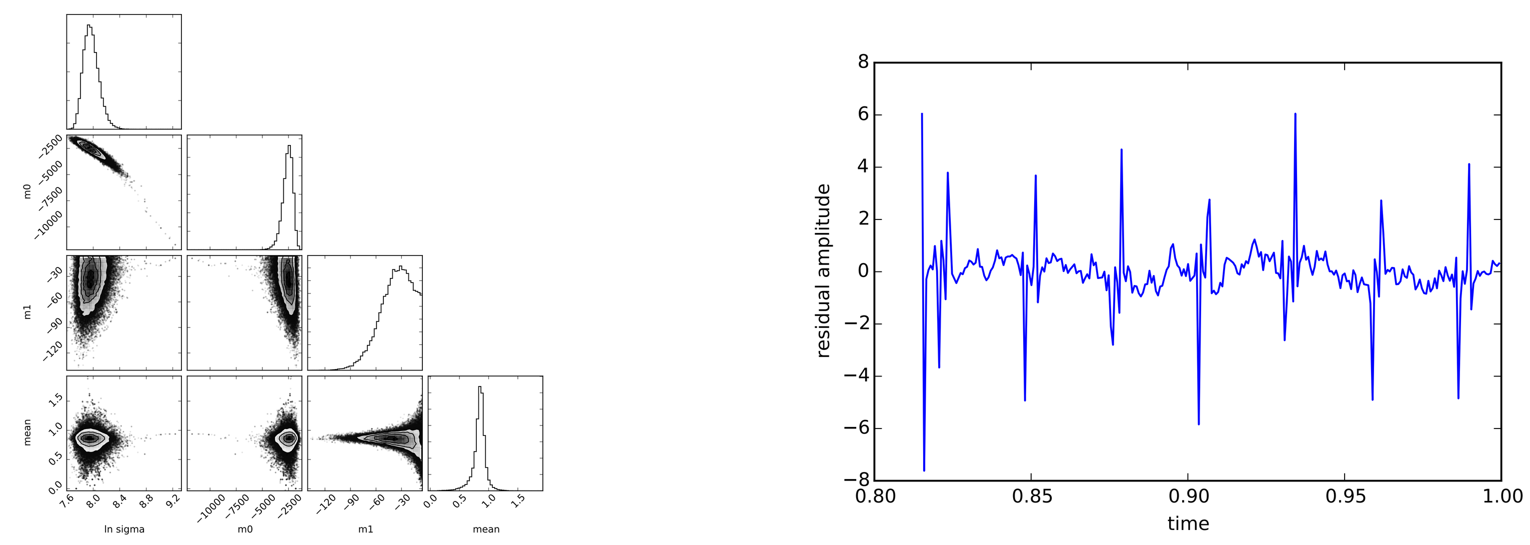

(a) (b)

Figure: (a) A 'triangle plot' showing the marginalised distributions of each combination of parameters. The 1D Gaussians give a distribution of the most likely value for each parameter $m_i = \frac{1}{r_i}$. (b) The signal that remains when all correlations have been removed from the noise. The transits can now clearly be seen, although their shape has been changed by the whitening filter.

Now that the data has been reduced to a form such that the remaining noise is $N(0,1)$, case specific analysis techniques can be used to explore the residual signal.

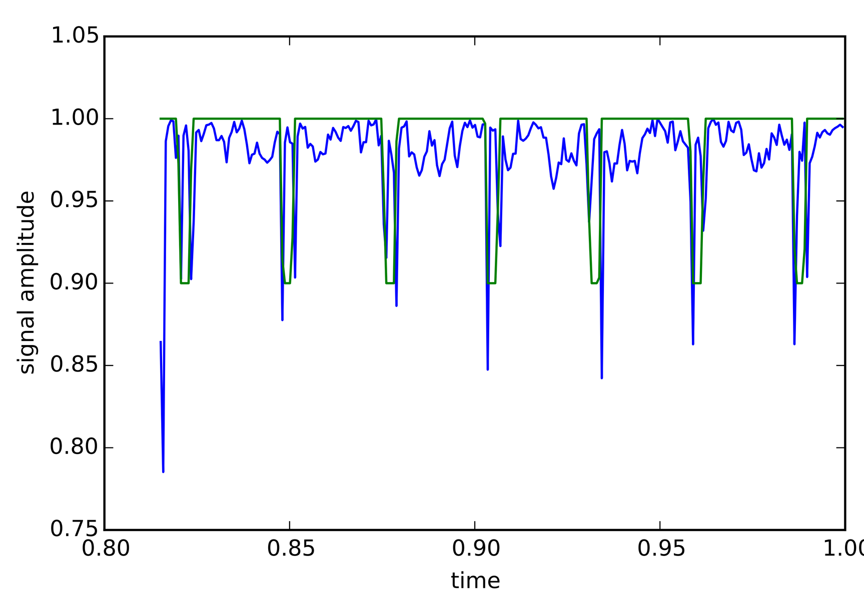

Figure: A perfect transit signal (green) with the fitted and manipulated noisy signal (blue).

More Information

This poster, and regular updates on this project can be found at `http://www.jimbarrett.co.uk`