Modelling coloured noise

Christian Röver

May 8th 2008

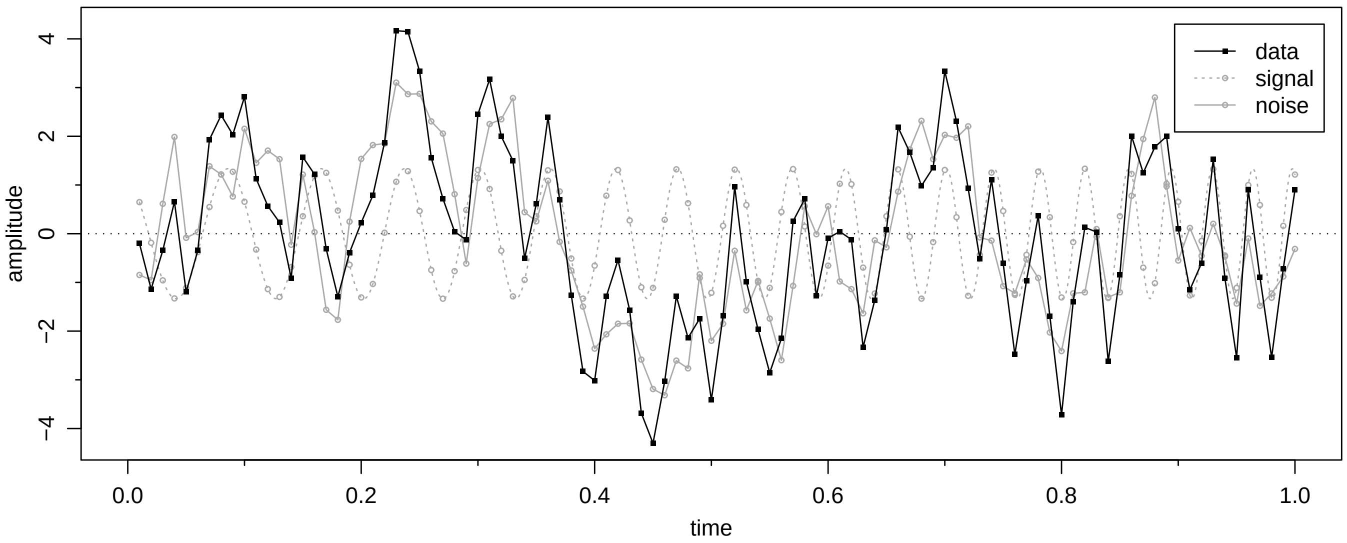

Measuring noisy signals

- common assumption:

$$\underbrace{y(t_i)}_{\text{data}} = \underbrace{s_{\vec{\vartheta}}(t_i)}_{\text{signal}} + \underbrace{n(t_i)}_{\text{noise}}$$

  where noise is assumed **gaussian** with **known power spectral density** S_n

- **technically**:
  Likelihood $p(y|\vec{\vartheta})$ based on Fourier transform of $y(t_i) - s_{\vec{\vartheta}}(t_i)$ . . .

$$\log\bigl(p(y|\vec{\vartheta})\bigr) \;\propto\; -\sum_j \frac{|\tilde{y}(f_j) - \tilde{s}_{\vec{\vartheta}}(f_j)|^2}{2\, S_n(f_j)}$$

- problem: noise spectrum unknown

- e.g. LISA: spectrum known to be **coloured** and **not smooth**,
  and: little data, which must be used to infer parameters **and** spectrum
  ("nuisance parameter").

- consider frequency-domain expression of noise:

$$n(t_i) \;=\; \frac{1}{\sqrt{N\Delta_t}} \sum_j a_j \cos(2\pi f_j t_i) + b_j \sin(2\pi f_j t_i)$$

$n(t_i)$ random $\quad\leftrightarrow\quad$ a_j, b_j random. . .

$\Rightarrow \quad \mathrm{E}[n(t_i)] = 0 \quad \Leftrightarrow \quad \mathrm{E}[a_j] = \mathrm{E}[b_j] = 0$

$\Rightarrow \quad \frac{\mathrm{Var}(a_j)+\mathrm{Var}(b_j)}{2} = S_n(f_j)$

Setting up the noise model

- **assume**: $\mathrm{E}[n(t_i)] = 0$ and $S_n(f_j) < \infty$

  $\Rightarrow$ **Maximum Entropy** distribution of a_j and b_j (for given S_n):

$$a_j, b_j \;\sim\; \mathrm{N}(\mu = 0,\, \sigma_j^2 = S_n(f_j))$$

  $\Rightarrow$ **likelihood**:

$$\log\bigl(p(n|\sigma_0^2, \ldots, \sigma_{N/2}^2)\bigr) \;\propto\; -\sum_j \left(\log(\sigma_j^2) + \frac{|\tilde{n}(f_j)|^2}{2\,\sigma_j^2}\right)$$

- **prior** distribution for σ_j^2: conjugate **scaled inverse** χ^2**-distribution** (convenient: posterior of same family; sensible: includes Jeffreys prior)

Implementation in usual "signal plus noise" framework

- model:

$$y(t_i) \;=\; s_{\vec{\vartheta}}(t_i) \,+\, n_{\vec{\sigma}}(t_i)$$

- likelihood:

$$\log\bigl(p(y|\vec{\vartheta},\vec{\sigma})\bigr) \;\propto\; -\sum_j \left(\log(\sigma_j^2) + \frac{|\tilde{y}(f_j) - \tilde{s}_{\vec{\vartheta}}(f_j)|^2}{2\,\sigma_j^2}\right)$$

- implementation: **Gibbs sampler**,
  alternatingly drawing from conditionals $p(\vec{\vartheta}|\vec{\sigma},y)$ and $p(\vec{\sigma}|\vec{\vartheta},y)$
  - $p(\vec{\vartheta}|\vec{\sigma},y)$: "as usual" (Metropolis- / Metropolis-Hastings)
  - $p(\vec{\sigma}|\vec{\vartheta},y)$: straightforward to draw from conditional Inv-$\chi^2(\cdot,\cdot)$ distn.

- **Example**: non-gaussian, coloured noise + chirping signal $s_{f,\dot{f},a,\phi}(t_i)$ (amplitude a, frequency f, frequency derivative $\dot{f}$, phase ϕ)