Bayesian Inference in the Local Level model

Hedibert Freitas Lopes

March 2014

Local level model

The first order Gaussian dynamic linear model (aka, local level model) is a basic measurement error model where the underlying (hidden) signal evolves according to a simple random walk process.

More precisely,

$$\begin{aligned} y_t|\beta_t,\sigma^2 &\sim N(\beta_t,\sigma^2) \\ \beta_t|\beta_{t-1},\tau^2 &\sim N(\beta_{t-1},\tau^2) \end{aligned}$$

for $t = 1, \ldots, n$.

Let $\theta = (\sigma^2, \tau^2, \beta_0)$.

Prior of θ

Let the prior of θ be conditionally conjugated, i.e.

$$\begin{aligned} p(\sigma^2, \tau^2, \beta_0) &= p_{IG}(\sigma^2; \nu_0/2, \nu_0\sigma_0^2/2) \\ &\times p_{IG}(\tau^2; v_0/2, v_0\tau_0^2/2) \\ &\times p_N(\beta_0; b_0, B_0) \end{aligned}$$

for known hyperparameters $(\nu_0, \sigma_0^2, v_0, \tau_0^2, b_0, B_0)$.

In our simulated and real data examples we set

$$\begin{aligned} \nu_0 &= 0.0001 \\ \sigma_0^2 &= 0.0001 \\ v_0 &= 0.0001 \\ \tau_0^2 &= 0.0001 \\ b_0 &= 0 \\ B_0 &= 10000 \end{aligned}$$

Full conditional posteriors of σ^2, τ^2 and β_0

It can be shown that

$$\begin{array}{rcl} \sigma^2|y,\beta & \sim & IG(\nu_1/2, \nu_1\sigma_1^2/2) \\ \tau^2|\beta,\beta_0 & \sim & IG(v_1/2, v_1\tau_1^2/2) \\ \beta_0|\beta_1,\tau^2 & \sim & N(b_1, B_1) \end{array}$$

where $\nu_1 = \nu_0 + n$, $v_1 = v_0 + n$,

$$\begin{array}{rcl} \nu_1\sigma_1^2 & = & \nu_0\sigma_0^2 + (y-\beta)'(y-\beta) \\ v_1\tau_1^2 & = & v_0\tau_0^2 + (\beta-\beta_{-1})'(\beta-\beta_{-1}) \\ B_1^{-1} & = & B_0^{-1} + \tau^{-2} \\ B_1^{-1}b_1 & = & B_0^{-1}b_0 + \tau^{-2}\beta_1 \end{array}$$

where $\beta_{-1} = (\beta_0, \beta_1, \ldots, \beta_{n-1})'$.

Prior of β

It can be easily seen that, conditionally on θ, the dynamics of β_t can be jointly described by the following multivariate normal distribution

$$\beta|\theta \sim N(\beta_0 1_n, \tau^2 A)$$

where

$$A = \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 1 & 2 & 2 & \cdots & 2 & 2 & 2 \\ 1 & 2 & 3 & \cdots & 3 & 3 & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 1 & 2 & 3 & \cdots & n-2 & n-2 & n-2 \\ 1 & 2 & 3 & \cdots & n-2 & n-1 & n-1 \\ 1 & 2 & 3 & \cdots & n-2 & n-1 & n \end{pmatrix}$$

Full conditional posterior of β

Combining likelihood, $y|\beta,\theta \sim N(\beta, \sigma^2 I_n)$, with prior, $\beta|\theta \sim N(\beta_0 1_n, \tau^2 A)$, leads to posterior

$$\beta|y,\theta \sim N\left\{(BA^{-1})\beta_0 + (\rho B)y, \tau^2 B\right\},$$

where $\rho = \tau^2/\sigma^2$, $B^{-1} = A^{-1} + \rho I_n$ and

$$B^{-1} = \begin{pmatrix} 2+\rho & -1 & 0 & \cdots & 0 & 0 & 0 \\ -1 & 2+\rho & -1 & \cdots & 0 & 0 & 0 \\ 0 & -1 & 2+\rho & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 2+\rho & -1 & 0 \\ 0 & 0 & 0 & \cdots & -1 & 2+\rho & -1 \\ 0 & 0 & 0 & \cdots & 0 & -1 & 2+\rho \end{pmatrix}.$$

Computational issue: Inverting A can be quite expensive!

Solution: forward filtering backward sampling

Forward filter (Kalman filter)

For $t = 0$: $(\beta_0|y^0) \sim N(m_0, C_0)$

For $t = 1, \ldots, n$:

$$\begin{array}{rcll} \beta_t|y^{t-1}, \theta & \sim & N(m_{t-1}, R_t) & \quad R_t = C_{t-1} + \tau^2 \\ y_t|y^{t-1}, \theta & \sim & N(m_{t-1}, Q_t) & \quad Q_t = C_{t-1} + \tau^2 + \sigma^2 \\ \beta_t|y^t, \theta & \sim & N(m_t, C_t) & \end{array}$$

where $m_t = (1 - A_t)m_{t-1} + A_t y_t$, $C_t = (1 - A_t)R_t$ and $A_t = Q_t/R_t$.

Backward sampler (Kalman smoother)

In addition, for $t = n - 1, n - 2, \ldots, 1$

$$\beta_t|\beta_{t+1}, \theta \sim N((1 - B_t)m_t + B_t\beta_{t+1}, (1 - B_t^2)C_t - B_t^2\tau^2)$$

where $B_t = C_t/(C_t + \tau^2)$.

Simulating $n = 200$: $(\sigma, \tau, \beta_0) = (1.0, 0.5, 0.0)$

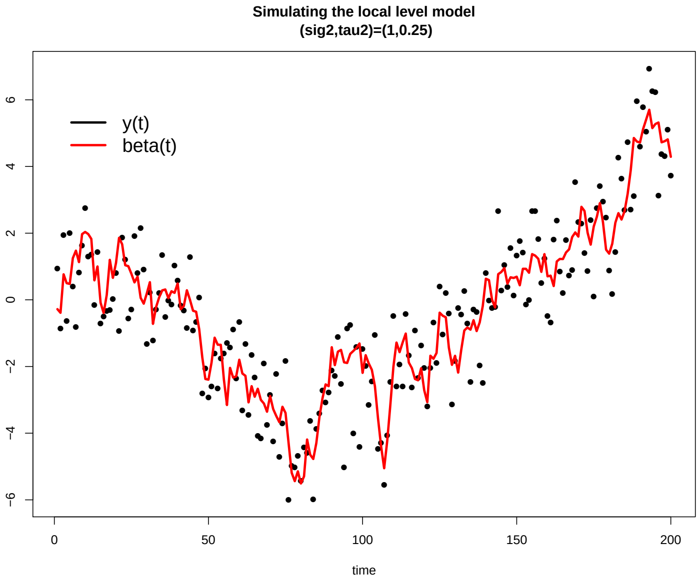

Bayesian linear regression versus FFBS

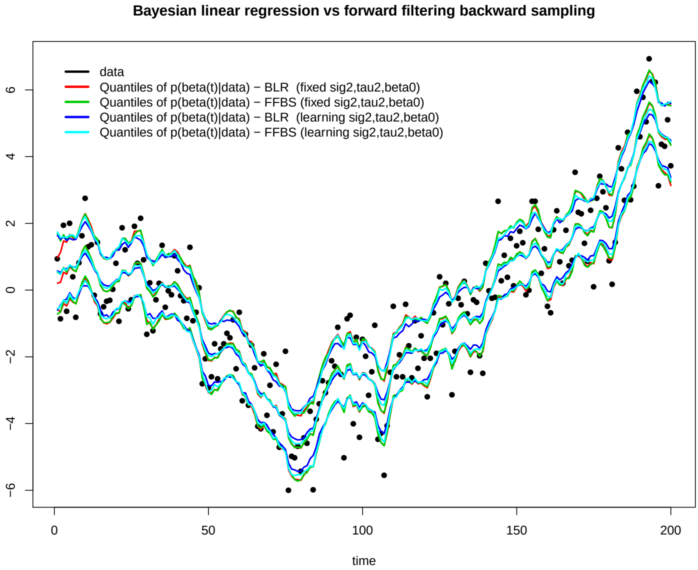

Running time