

CPSC 535

Gibbs Sampling

AD

February 2007

Finite State-Space Hidden Markov Models

- Mixture models cannot model dependent data; one straightforward extension consists of picking for $\{X_n\}$ a finite state-space Markov chain.

Finite State-Space Hidden Markov Models

- Mixture models cannot model dependent data; one straightforward extension consists of picking for $\{X_n\}$ a finite state-space Markov chain.
- We have say $X_n \in \{1, \dots, K\}$ with

$$Y_n | X_n \sim g_{X_n}(y)$$

but $\Pr(X_1 = i) = \mu_i$ and

$$\Pr(X_{n+1} = j | X_n = i) = p_{i,j}.$$

Finite State-Space Hidden Markov Models

- Mixture models cannot model dependent data; one straightforward extension consists of picking for $\{X_n\}$ a finite state-space Markov chain.
- We have say $X_n \in \{1, \dots, K\}$ with

$$Y_n | X_n \sim g_{X_n}(y)$$

but $\Pr(X_1 = i) = \mu_i$ and

$$\Pr(X_{n+1} = j | X_n = i) = p_{i,j}.$$

- In this case, the probability to stay in a given state is geometric.

Finite State-Space Hidden Markov Models

- Mixture models cannot model dependent data; one straightforward extension consists of picking for $\{X_n\}$ a finite state-space Markov chain.
- We have say $X_n \in \{1, \dots, K\}$ with

$$Y_n | X_n \sim g_{X_n}(y)$$

but $\Pr(X_1 = i) = \mu_i$ and

$$\Pr(X_{n+1} = j | X_n = i) = p_{i,j}.$$

- In this case, the probability to stay in a given state is geometric.
- Simple model (over)used in speech processing, DNA sequence analysis, communications etc.




Figure: Realization of 100 observations for $K = 3$, $\mu_1 = -1$, $\sigma_1^2 = 0.1$, $\mu_2 = 0$, $\sigma_2^2 = 1$, $\mu_3 = 1$, $\sigma_3^2 = 0.1$ with $p_{i,i} = 0.90$, $p_{i,j} = 0.05$ for $i \neq j$. $\{X_n\}$ is displayed in red, $\{Y_n\}$ in blue.

- Given T observations y_1, \dots, y_T then the likelihood of the observations is given by

$$p(y_1, \dots, y_T | \theta)$$

where θ includes all the unknown parameters.

- Given T observations y_1, \dots, y_T then the likelihood of the observations is given by

$$p(y_1, \dots, y_T | \theta)$$

where θ includes all the unknown parameters.

- The likelihood can be computed exactly using a simple recursion. However, we limit ourselves first to the complete likelihood

$$p(y_{1:T}, x_{1:T} | \theta) = p(y_{1:T} | \theta, x_{1:T}) p(x_{1:T} | \theta)$$

where

$$\begin{aligned} p(y_{1:T} | \theta, x_{1:T}) &= \prod_{n=1}^T p(y_n | \theta, x_n), \\ p(x_{1:T} | \theta) &= p(x_1 | \theta) \prod_{n=2}^T p(x_n | \theta, x_{n-1}). \end{aligned}$$

- Typically, one uses the EM algorithm to estimate the maximum likelihood estimate of the unknown parameter θ .

- Typically, one uses the EM algorithm to estimate the maximum likelihood estimate of the unknown parameter θ .
- Alternatively, given a prior distribution $p(\theta)$ on θ , then we can perform Bayesian inference and estimate

$$p(\theta, x_{1:T} | y_{1:T}) = \frac{p(y_{1:T} | \theta, x_{1:T}) p(x_{1:T} | \theta) p(\theta)}{p(y_{1:T})}$$

- Typically, one uses the EM algorithm to estimate the maximum likelihood estimate of the unknown parameter θ .
- Alternatively, given a prior distribution $p(\theta)$ on θ , then we can perform Bayesian inference and estimate

$$p(\theta, x_{1:T} | y_{1:T}) = \frac{p(y_{1:T} | \theta, x_{1:T}) p(x_{1:T} | \theta) p(\theta)}{p(y_{1:T})}$$

- For mixture, there is no closed-form. Hence there is none for HMM. The Gibbs sampler can be implemented for this class of models by sampling iteratively from $p(\theta | y_{1:T}, x_{1:T})$ and $p(x_{1:T} | y_{1:T}, \theta)$.

Extension to General State-Space HMM

- It is important to realize that this class of models can be significantly extended by taking a latent process $\{X_n\}$ which is not discrete-valued.

Extension to General State-Space HMM

- It is important to realize that this class of models can be significantly extended by taking a latent process $\{X_n\}$ which is not discrete-valued.
- A simple example correspond to the case where

$$X_n = \alpha X_{n-1} + \sigma_v V_n, \quad V_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

$$Y_n = X_n + \sigma_w W_n, \quad W_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

Extension to General State-Space HMM

- It is important to realize that this class of models can be significantly extended by taking a latent process $\{X_n\}$ which is not discrete-valued.
- A simple example correspond to the case where

$$X_n = \alpha X_{n-1} + \sigma_v V_n, \quad V_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

$$Y_n = X_n + \sigma_w W_n, \quad W_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

- Clearly, we are in the case where $\{X_n\}$ is a Markov process

$$X_n | X_{n-1} \sim f_\theta(x_n | x_{n-1})$$

and $Y_n | X_n \sim g_\theta(y_n | x_n)$ where

$$f_\theta(x_n | x_{n-1}) = \mathcal{N}(x_n; \alpha x_{n-1}, \sigma_v^2),$$

$$g_\theta(y_n | x_n) = \mathcal{N}(y_n; x_n, \sigma_w^2).$$

and $\theta = (\alpha, \sigma_v^2, \sigma_w^2)$.

- Suppose you have

$$Y_n = g(t_n) + W_n \text{ where } W_n \sim \mathcal{N}(0, \sigma^2)$$

with

$$\frac{d^2 g(t)}{dt^2} = \tau \frac{dB(t)}{dt} \text{ where } B(t) \text{ Wiener process}$$

with $B(0) = 0$ and $\text{var}(B(t)) = 1$.

- Suppose you have

$$Y_n = g(t_n) + W_n \text{ where } W_n \sim \mathcal{N}(0, \sigma^2)$$

with

$$\frac{d^2 g(t)}{dt^2} = \tau \frac{dB(t)}{dt} \text{ where } B(t) \text{ Wiener process}$$

with $B(0) = 0$ and $\text{var}(B(t)) = t$.

- With initial conditions such that $(g(t_1) \quad dg(t_1)/dt) \sim \mathcal{N}(0, kl_2)$

$$Y_n = (1 \ 0) X(t_n) + W_n,$$

$$X(t_n) = \begin{pmatrix} 1 & \delta_n \\ 0 & 1 \end{pmatrix} X(t_{n-1}) + V_n, \quad V_n \sim \mathcal{N}\left(0, \begin{pmatrix} \delta_n^3/3 & \delta_n^2/2 \\ \delta_n^2/2 & \delta_n \end{pmatrix}\right)$$

where $\delta_n = t_n - t_{n-1}$.




Figure: Bearings-only-tracking data

- Consider the coordinates of a target observed through a radar.

$$\begin{pmatrix} X_n^1 \\ \cdot 1 \\ X_n \\ X_n^2 \\ \cdot 2 \\ X_n \end{pmatrix} = \Delta \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} X_{n-1}^1 \\ \cdot 1 \\ X_{n-1} \\ X_{n-1}^2 \\ \cdot 2 \\ X_{n-1} \end{pmatrix} + V_n, \quad V_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma_v),$$

$$Y_n = \tan^{-1} \left(\frac{X_n^1}{X_n^2} \right) + W_n, \quad W_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2).$$

where the process $\{Y_n\}$ is observed but $\{X_n\}$ is unknown.




Figure: Four stock prices




Figure: Log-return of a stock price

- Consider the log-return sequence of a stock then a popular model in financial econometrics is the stochastic volatility model

$$X_n = \alpha X_{n-1} + \sigma V_n \text{ where } V_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

$$Y_n = \beta \exp(X_n/2) W_n \text{ where } W_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

where the process $\{Y_n\}$ is observed but $\{X_n\}$ and $\theta = (\alpha, \sigma, \beta)$ are unknown.

- Consider the log-return sequence of a stock then a popular model in financial econometrics is the stochastic volatility model

$$X_n = \alpha X_{n-1} + \sigma V_n \text{ where } V_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

$$Y_n = \beta \exp(X_n/2) W_n \text{ where } W_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

where the process $\{Y_n\}$ is observed but $\{X_n\}$ and $\theta = (\alpha, \sigma, \beta)$ are unknown.

- We have

$$f_{\theta}(x_n | x_{n-1}) = \mathcal{N}(x_n; \alpha x_{n-1}, \sigma_v^2),$$

$$g_{\theta}(y_n | x_n) = \mathcal{N}(y_n; 0, \beta^2 \exp(x_n)).$$

- Many real-world problems can be rewritten as

$$\begin{aligned} X_n | X_{n-1} &\sim f_\theta(x_n | x_{n-1}), \quad X_1 \sim \mu(x_1), \\ Y_n | X_n &\sim g_\theta(y_n | x_n) \end{aligned}$$

where $\theta \sim p(\theta)$.

- Many real-world problems can be rewritten as

$$\begin{aligned} X_n | X_{n-1} &\sim f_\theta(x_n | x_{n-1}), \quad X_1 \sim \mu(x_1), \\ Y_n | X_n &\sim g_\theta(y_n | x_n) \end{aligned}$$

where $\theta \sim p(\theta)$.

- In a Bayesian framework, given $y_{1:T}$, we are interested in estimating the posterior

$$p(x_{1:T}, \theta | y_{1:T}) \propto p(y_{1:T} | \theta, x_{1:T}) p(x_{1:T} | \theta) p(\theta)$$

where

$$\begin{aligned} p(y_{1:T} | \theta, x_{1:T}) &= \prod_{n=1}^T g_\theta(y_n | x_n), \\ p(x_{1:T} | \theta) &= \mu(x_1) \prod_{n=2}^T f_\theta(x_n | x_{n-1}). \end{aligned}$$

- Assume you have

$$X_n = \alpha X_{n-1} + \sigma_v V_n, \quad V_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

$$Y_n = X_n + \sigma_w W_n, \quad W_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

where $X_1 \sim \mathcal{N}(0, 1)$, $\alpha \sim \mathcal{N}(0, \sigma_0^2)$, $\sigma_v^2 \sim \text{IG}(\frac{\nu_0}{2}, \frac{\gamma_0}{2})$ and $\sigma_w^2 \sim \text{IG}(\frac{\nu_0}{2}, \frac{\gamma_0}{2})$.

- Assume you have

$$X_n = \alpha X_{n-1} + \sigma_v V_n, \quad V_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

$$Y_n = X_n + \sigma_w W_n, \quad W_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

where $X_1 \sim \mathcal{N}(0, 1)$, $\alpha \sim \mathcal{N}(0, \sigma_0^2)$, $\sigma_v^2 \sim \text{IG}(\frac{\nu_0}{2}, \frac{\gamma_0}{2})$ and $\sigma_w^2 \sim \text{IG}(\frac{\nu_0}{2}, \frac{\gamma_0}{2})$.

- Gibbs sampler based on

$$p(x_k | y_{1:T}, x_{-k}, \alpha, \sigma_v^2, \sigma_w^2), \quad p(\sigma_v^2, \sigma_w^2 | y_{1:T}, x_{1:T}, \alpha), \\ p(\alpha | y_{1:T}, x_{1:T}, \sigma_v^2, \sigma_w^2).$$

- We have for $1 < k < T$

$$\begin{aligned}
 p(x_k | y_{1:T}, x_{-k}, \alpha, \sigma_v^2, \sigma_w^2) &\propto g(y_k | x_k, \sigma_w^2) f(x_k | x_{k-1}, \alpha, \sigma_v^2) \\
 &\quad \times f(x_{k+1} | x_k, \alpha, \sigma_v^2) \\
 &= \mathcal{N}(x_k; m_k, \sigma_k^2)
 \end{aligned}$$

where

$$\begin{aligned}
 m_k &= \sigma_k^2 \left(\frac{y_k^2}{\sigma_k^2} + \alpha \frac{x_{k+1} + x_{k-1}}{\sigma_v^2} \right), \\
 \frac{1}{\sigma_k^2} &= \frac{1}{\sigma_w^2} + \frac{\alpha^2 + 1}{\sigma_v^2}.
 \end{aligned}$$

- We have for $1 < k < T$

$$\begin{aligned}
 p(x_k | y_{1:T}, x_{-k}, \alpha, \sigma_v^2, \sigma_w^2) &\propto g(y_k | x_k, \sigma_w^2) f(x_k | x_{k-1}, \alpha, \sigma_v^2) \\
 &\quad \times f(x_{k+1} | x_k, \alpha, \sigma_v^2) \\
 &= \mathcal{N}(x_k; m_k, \sigma_k^2)
 \end{aligned}$$

where

$$\begin{aligned}
 m_k &= \sigma_k^2 \left(\frac{y_k^2}{\sigma_k^2} + \alpha \frac{x_{k+1} + x_{k-1}}{\sigma_v^2} \right), \\
 \frac{1}{\sigma_k^2} &= \frac{1}{\sigma_w^2} + \frac{\alpha^2 + 1}{\sigma_v^2}.
 \end{aligned}$$

- We have

$$p(\sigma_v^2, \sigma_w^2 | y_{1:T}, x_{1:T}, \alpha) = p(\sigma_v^2 | x_{1:T}, \alpha) p(\sigma_w^2 | y_{1:T}, x_{1:T})$$

- We have

$$\begin{aligned} p(\sigma_v^2 | x_{1:T}, \alpha) &\propto p(x_{1:T} | \alpha, \sigma_v^2) p(\sigma_v^2) \\ &\propto \frac{1}{\sigma_v^{T-1}} \exp\left(-\frac{\sum_{k=2}^T (x_k - \alpha x_{k-1})^2}{2\sigma_v^2}\right) \frac{1}{\sigma_v^{\nu_0}} \exp\left(-\frac{\gamma_0}{2\sigma_v^2}\right) \\ &= \mathcal{IG}\left(\sigma_v^2; \frac{\nu_0 + T - 1}{2}, \frac{\gamma_0 + \sum_{k=2}^T (x_k - \alpha x_{k-1})^2}{2}\right) \end{aligned}$$

- We have

$$\begin{aligned}
 p(\sigma_v^2 | x_{1:T}, \alpha) &\propto p(x_{1:T} | \alpha, \sigma_v^2) p(\sigma_v^2) \\
 &\propto \frac{1}{\sigma_v^{T-1}} \exp\left(-\frac{\sum_{k=2}^T (x_k - \alpha x_{k-1})^2}{2\sigma_v^2}\right) \frac{1}{\sigma_v^{v_0}} \exp\left(-\frac{\gamma_0}{2\sigma_v^2}\right) \\
 &= \mathcal{IG}\left(\sigma_v^2; \frac{v_0+T-1}{2}, \frac{\gamma_0 + \sum_{k=2}^T (x_k - \alpha x_{k-1})^2}{2}\right)
 \end{aligned}$$

- We have

$$\begin{aligned}
 p(\sigma_w^2 | y_{1:T}, x_{1:T}) &\propto p(y_{1:T} | x_{1:T}, \sigma_w^2) p(\sigma_w^2) \\
 &\propto \frac{1}{\sigma_w^T} \exp\left(-\frac{\sum_{k=2}^T (y_k - x_k)^2}{2\sigma_w^2}\right) \frac{1}{\sigma_w^{v_0}} \exp\left(-\frac{\gamma_0}{2\sigma_w^2}\right) \\
 &= \mathcal{IG}\left(\sigma_w^2; \frac{v_0+T}{2}, \frac{\gamma_0 + \sum_{k=1}^T (y_k - x_k)^2}{2}\right)
 \end{aligned}$$

- Finally we have

$$\begin{aligned}
 p(\alpha | y_{1:T}, x_{1:T}, \sigma_v^2, \sigma_w^2) &= p(\alpha | x_{1:T}, \sigma_v^2) \propto p(x_{1:T} | \alpha, \sigma_v^2) p(\alpha) \\
 &\propto \frac{1}{\sigma_v^{T-1}} \exp\left(-\frac{\sum_{k=2}^T (x_k - \alpha x_{k-1})^2}{2\sigma_v^2}\right) \exp\left(-\frac{\alpha^2}{2\sigma_0^2}\right) \\
 &= \mathcal{N}(\alpha; m_\alpha, \sigma_\alpha^2)
 \end{aligned}$$

where

$$\begin{aligned}
 \frac{1}{\sigma_\alpha^2} &= \frac{1}{\sigma_0^2} + \frac{\sum_{k=1}^{T-1} x_k^2}{\sigma_v^2}, \\
 m_\alpha &= \sigma_\alpha^2 \left(\sum_{k=2}^T x_k x_{k-1} \right).
 \end{aligned}$$




Figure: 100,000 samples after 10,000 burn in with $\alpha = 0.9$, $\sigma_w = 1$ and $\sigma_v = 1$ for $T = 100$. Approximations of $p(\alpha | y_{1:T})$, $p(\sigma_w^2 | y_{1:T})$ and $p(\sigma_v^2 | y_{1:T})$

- We have

$$X_n = AX_{n-1} + V_n, \quad V_n \sim \mathcal{N}(0, \Sigma),$$

$$Y_n = \tan^{-1} \left(\frac{X_n^1}{X_n^2} \right) + W_n, \quad W_n \sim \mathcal{N}(0, \sigma^2)$$

- We have

$$X_n = AX_{n-1} + V_n, \quad V_n \sim \mathcal{N}(0, \Sigma),$$

$$Y_n = \tan^{-1} \left(\frac{X_n^1}{X_n^2} \right) + W_n, \quad W_n \sim \mathcal{N}(0, \sigma^2)$$

- Assume for sake of simplicity that only $x_{1:T}$ are unknown, we want to estimate

$$p(x_{1:T} | y_{1:T}).$$

- We sample from the full conditional distributions

$$\begin{aligned} p(x_k | y_{1:T}, x_{-k}) &\propto p(x_k | x_{-k}) g(y_k | x_k) \\ &\propto f(x_{k+1} | x_k) f(x_k | x_{k-1}) g(y_k | x_k). \end{aligned}$$

- We sample from the full conditional distributions

$$\begin{aligned} p(x_k | y_{1:T}, x_{-k}) &\propto p(x_k | x_{-k}) g(y_k | x_k) \\ &\propto f(x_{k+1} | x_k) f(x_k | x_{k-1}) g(y_k | x_k). \end{aligned}$$

- We have

$$p(x_k | x_{-k}) \propto f(x_{k+1} | x_k) f(x_k | x_{k-1}) = \mathcal{N}(x_k; m_k, \Sigma_k)$$

where

$$\begin{aligned} \Sigma_k^{-1} &= \Sigma^{-1} + A^T \Sigma^{-1} A, \\ m_k &= \Sigma_k (\Sigma^{-1} A x_{k-1} + A^T \Sigma^{-1} x_{k+1}) \end{aligned}$$

- To sample from

$$p(x_k | y_{1:T}, x_{-k}) \propto p(x_k | x_{-k}) g(y_k | x_k)$$

we can use rejection sampling as you can sample from $p(x_k | x_{-k})$ and

$$\begin{aligned} g(y_k | x_k) &= \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\left(y_k - \tan^{-1}\left(\frac{x_k^1}{x_k^2}\right)\right)^2 / (2\sigma^2)\right) \\ &\leq \frac{1}{\sqrt{2\pi}\sigma}. \end{aligned}$$

- To sample from

$$p(x_k | y_{1:T}, x_{-k}) \propto p(x_k | x_{-k}) g(y_k | x_k)$$

we can use rejection sampling as you can sample from $p(x_k | x_{-k})$ and

$$\begin{aligned} g(y_k | x_k) &= \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\left(y_k - \tan^{-1}\left(\frac{x_k^1}{x_k^2}\right)\right)^2 / (2\sigma^2)\right) \\ &\leq \frac{1}{\sqrt{2\pi}\sigma}. \end{aligned}$$

- Gibbs sampling can be implemented even for non-linear models




Figure: MCMC for state estimation using bearings-only-tracking data. Mean and credible intervals for $p(x_n | Y_{1:n})$.

- We have

$$X_n = \alpha X_{n-1} + \sigma V_n \text{ where } V_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

$$Y_n = \beta \exp(X_n/2) W_n \text{ where } W_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

- We have

$$X_n = \alpha X_{n-1} + \sigma V_n \text{ where } V_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

$$Y_n = \beta \exp(X_n/2) W_n \text{ where } W_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$$

- Prior model: $\alpha \sim \mathcal{U}(-1, 1)$, $\sigma^2 \sim \mathcal{IG}(\frac{\nu_0}{2}, \frac{\gamma_0}{2})$ and $\beta \sim \mathcal{IG}(\frac{\nu_0}{2}, \frac{\gamma_0}{2})$.

- We want to sample from

$$p(x_k | x_{-k}, y_{1:T}, \alpha, \sigma^2, \beta) \propto f(x_k | x_{k-1}, \alpha, \sigma^2) \\ \times f(x_{k+1} | x_k, \alpha, \sigma^2) g(y_k | x_k, \beta)$$

where

$$p(x_k | x_{-k}, \alpha, \sigma^2) \propto f(x_k | x_{k-1}, \alpha, \sigma^2) f(x_{k+1} | x_k, \alpha, \sigma^2) \\ = \mathcal{N}\left(x_k; m_k = \frac{\alpha(x_{k-1} + x_{k+1})}{1 + \alpha^2}, \sigma_k^2 = \frac{\sigma^2}{1 + \alpha^2}\right).$$

- We want to sample from

$$p(x_k | x_{-k}, y_{1:T}, \alpha, \sigma^2, \beta) \propto f(x_k | x_{k-1}, \alpha, \sigma^2) \times f(x_{k+1} | x_k, \alpha, \sigma^2) g(y_k | x_k, \beta)$$

where

$$\begin{aligned} p(x_k | x_{-k}, \alpha, \sigma^2) &\propto f(x_k | x_{k-1}, \alpha, \sigma^2) f(x_{k+1} | x_k, \alpha, \sigma^2) \\ &= \mathcal{N}\left(x_k; m_k = \frac{\alpha(x_{k-1} + x_{k+1})}{1 + \alpha^2}, \sigma_k^2 = \frac{\sigma^2}{1 + \alpha^2}\right). \end{aligned}$$

- We have

$$\begin{aligned} \log g(y_k | x_k, \beta) &\equiv -\frac{x_k}{2} - \frac{y_k^2}{2\beta^2} \exp(-x_k) \\ &\leq -\frac{x_k}{2} - \frac{y_k^2}{2\beta^2} (\exp(-m_k)(1 + m_k) - x_k \exp(-m_k)) \quad [\text{as } \exp(u) \geq 1 + u] \\ &= \log g^*(y_k | x_k, \beta) \end{aligned}$$

- We propose to sample from $p(x_k | x_{k-1}, x_{k+1}, y_k, \alpha, \sigma^2, \beta)$ using rejection by sampling from where

$$\begin{aligned} q(x_k) &\propto p(x_k | x_{-k}, \alpha, \sigma^2) g^*(y_k | x_k, \beta) \\ &= \mathcal{N}\left(x_k; m_k + \frac{\sigma_k^2}{2} \left[\frac{y_k^2}{\beta_2} \exp(-m_k^2) - 1 \right], \sigma_k^2\right). \end{aligned}$$

- We propose to sample from $p(x_k | x_{k-1}, x_{k+1}, y_k, \alpha, \sigma^2, \beta)$ using rejection by sampling from where

$$\begin{aligned}
 q(x_k) &\propto p(x_k | x_{-k}, \alpha, \sigma^2) g^*(y_k | x_k, \beta) \\
 &= \mathcal{N}\left(x_k; m_k + \frac{\sigma_k^2}{2} \left[\frac{y_k^2}{\beta_2} \exp(-m_k^2) - 1 \right], \sigma_k^2\right).
 \end{aligned}$$

- Then we accept the proposal with probability

$$\frac{g(y_k | x_k, \beta)}{g^*(y_k | x_k, \beta)}.$$

- We propose to sample from $p(x_k | x_{k-1}, x_{k+1}, y_k, \alpha, \sigma^2, \beta)$ using rejection by sampling from where

$$\begin{aligned}
 q(x_k) &\propto p(x_k | x_{-k}, \alpha, \sigma^2) g^*(y_k | x_k, \beta) \\
 &= \mathcal{N}\left(x_k; m_k + \frac{\sigma_k^2}{2} \left[\frac{y_k^2}{\beta_2} \exp(-m_k^2) - 1 \right], \sigma_k^2\right).
 \end{aligned}$$

- Then we accept the proposal with probability

$$\frac{g(y_k | x_k, \beta)}{g^*(y_k | x_k, \beta)}.$$

- Update of the hyperparameters are straightforward.




Figure: UK Sterling/US dollar exchange rates from 1/10/81 to 28/6/85: 200,000 samples after 20,000 burn-in. Approximations of $p(\alpha | y_{1:T})$, $p(\sigma^2 | y_{1:T})$ and $p(\beta | y_{1:T})$.

- These Gibbs sampling algorithms are simple but once more they are not very efficient as we sample typically $p(x_k | y_{1:T}, x_{-k}, \theta)$ then $p(\theta | y_{1:T}, x_{1:T})$.

- These Gibbs sampling algorithms are simple but once more they are not very efficient as we sample typically $p(x_k | y_{1:T}, x_{-k}, \theta)$ then $p(\theta | y_{1:T}, x_{1:T})$.
- We would like to be able to sample all the states variables jointly; i.e. sampling iteratively from $p(x_{1:T} | y_{1:T}, \theta)$ then $p(\theta | y_{1:T}, x_{1:T})$.

- These Gibbs sampling algorithms are simple but once more they are not very efficient as we sample typically $p(x_k | y_{1:T}, x_{-k}, \theta)$ then $p(\theta | y_{1:T}, x_{1:T})$.
- We would like to be able to sample all the states variables jointly; i.e. sampling iteratively from $p(x_{1:T} | y_{1:T}, \theta)$ then $p(\theta | y_{1:T}, x_{1:T})$.
- Generally sampling exactly from $p(x_{1:T} | y_{1:T}, \theta)$ is impossible except for HMM and linear Gaussian models.

- All the models we have seen rely on the ability to sample from some full conditional distribution $\pi(\theta_k | \theta_{-k})$.

- All the models we have seen rely on the ability to sample from some full conditional distribution $\pi(\theta_k | \theta_{-k})$.
- Although it is possible in numerous models, there are also numerous models where one CANNOT do it.

- All the models we have seen rely on the ability to sample from some full conditional distribution $\pi(\theta_k | \theta_{-k})$.
- Although it is possible in numerous models, there are also numerous models where one CANNOT do it.
- In such cases, alternative methods relying on the Metropolis-Hastings algorithm have to be developed.