

STAT 4010 – Bayesian Learning

TUTORIAL 9

Spring 2022

Cheuk Hin (Andy) CHENG ([Email](#) | [Homepage](#))

Di SU ([Email](#) | [Homepage](#))

1 Reweighting Importance Sampling

Recall that in importance sampling, we aim at simulating

$$I = \int_{\Theta} g(\theta)\pi_u(\theta)d\theta = \int_{\Theta} g(\theta)\frac{\pi_u(\theta)}{p(\theta)}p(\theta)d\theta = \int_{\Theta} g(\theta)w(\theta)p(\theta)d\theta,$$

where $w(\theta) = \pi_u(\theta)/p(\theta)$. We can use the reweighting importance sampling (RIS) to generate samples from target density.

Algorithm 1: Reweighting Importance Sampling.

Input: (i) simulation size J ; (ii) proposed PDF $p(\cdot)$; (iii) (unnormalized) target density $\pi_u(\cdot)$.

begin

(1) Generate $\tilde{\theta}_1, \dots, \tilde{\theta}_J \stackrel{iid}{\sim} p(\cdot)$.

(2) Compute $w_j = \pi_u(\tilde{\theta}_j)/p(\tilde{\theta}_j)$, for $j = 1, \dots, J$.

(3) **for** j in $\{1, \dots, J\}$ **do**

 Draw one element from $\{\tilde{\theta}_1, \dots, \tilde{\theta}_j\}$ with probabilities $\{w_1, \dots, w_j\}$.

 Saved the sampled element to θ_j

end

end

Output: $\theta_{1:J}$

Theorem 1.1. (Justification of the RIS) If $\text{Var}_p(w(\theta)) < \infty$, then as $J \rightarrow \infty$ for all $j = 1, \dots, J$ and $t \in \mathbb{R}$,

$$\Pr(\theta_j \leq t) \rightarrow \int_{-\infty}^t \pi(\theta)d\theta.$$

Remark 1.1. For estimating I using the RIS, we may consider

$$\hat{I}_{RIS} = \frac{1}{J} \sum_{j=1}^J g(\theta_j),$$

where $\{\theta_1, \dots, \theta_J\}$ are the outputs of RIS.

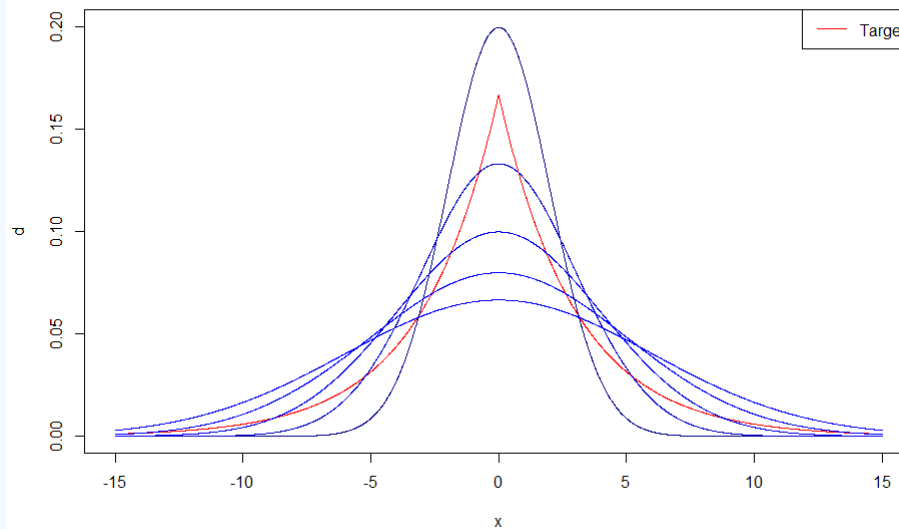
Example 1.1. Consider $\theta \sim F(\theta)$ and $f(\theta) \propto \exp\{-|\theta|/3\}\mathbb{1}(\theta \in \mathbb{R})$. Simulate $E\theta^2$ using RIS.

SOLUTION: Note that we use a Normal distribution as the proposal following last tutorial. Also, $E\theta^2 = 18$.

```

1 target_den <- function(theta,b=3){
2   log_d = -abs(theta)/b
3   exp(log_d - max(log_d))
4 }
5
6 target_g <- function(theta){
7   theta^2
8 }
9
10 ##RIS
11 RIS <- function(J,sd = 4, b= 3){ ##sd being the sd of the proposal
12   tilde_theta = rnorm(J,0,sd)
13   w = target_den(tilde_theta,b)/dnorm(tilde_theta,0,sd)
14   sample(tilde_theta,size = J,replace=T,prob = w) ##output the theta_{1;J}
15 }
16
17 sim_I_RIS <- function(J,sd=4,b=3){
18   theta_sim = RIS(J,sd,b)
19   mean(target_g(theta_sim))
20 }
21
22 J_sim = 2^14
23 sim_I_RIS(J_sim,sd = 2)
24 [1] 10.91392
25 sim_I_RIS(J_sim,sd = 3)
26 [1] 13.19787
27 sim_I_RIS(J_sim,sd = 4)
28 [1] 15.00232
29 sim_I_RIS(J_sim,sd = 5)
30 [1] 17.36904
31 sim_I_RIS(J_sim,sd = 6)
32 [1] 18.6439
33 sim_I_RIS(J_sim,sd = 7)
34 [1] 18.78481
35 sim_I_RIS(J_sim,sd = 8)
36 [1] 17.99314
37
38 x = seq(-15,15,0.01)
39 d = exp(-abs(x)/3)/(2*3)
40 plot(x,d,type='l',ylim=c(0,0.2),col='red',main = 'Target Density vs Proposal')
41 lines(x,dnorm(x,0,2),col='blue4')
42 lines(x,dnorm(x,0,3),col='blue3')
43 lines(x,dnorm(x,0,4),col='blue2')
44 lines(x,dnorm(x,0,5),col='blue1')
45 lines(x,dnorm(x,0,6),col='blue')
46 legend('topright',legend=c('Target'),col='red',lwd=1.5)

```



2 Markov Chain Monte Carlo (MCMC)

2.1 Theoretical Background

Definition 1. Let θ_j be a Θ -valued RV for each $j = 1, \dots, J$.

1. A collection of RVs $\{\theta_j : j = 1, 2, \dots\}$ is said to be a Markov chain if

$$\forall j \in \{1, \dots, J\}, \quad [\theta_j \mid \theta_{1:(j-1)}] = [\theta_j \mid \theta_{j-1}].$$

2. A Markov chain $\{\theta_j\}$ is said to be time-homogeneous (or simply homogenous) if

$$\forall \delta, \quad \forall k, \quad \forall j_0 \leq j_1 \leq \dots \leq j_k, \quad [\theta_{j_1}, \dots, \theta_{j_k} \mid \theta_{j_0}] = [\theta_{j_1+\delta}, \dots, \theta_{j_k+\delta} \mid \theta_{j_0+\delta}]$$

3. The transition kernel $K(\cdot \mid \cdot)$ of a homogeneous Markov chain $\{\theta_j\}$ is the conditional density of $[\theta_j \mid \theta_{j-1}]$, i.e.,

$$K(\theta_j \mid \theta_{j-1}) = f(\theta_j \mid \theta_{j-1}).$$

Theorem 2.1. (Important distribution related to a Markov Chain)

1. A Markov chain $\{\theta_j\}$ is said to have a stationary distribution π if

$$\pi(\theta) = \int_{\Theta} K(\theta \mid \theta') \pi(\theta') d\theta'.$$

2. A Markov chain $\{\theta_j\}$ is said to satisfy the detailed balance condition with a probability density π if

$$\forall \theta, \theta' \in \Theta, \quad K(\theta \mid \theta') \pi(\theta') = K(\theta' \mid \theta) \pi(\theta).$$

3. (Detailed balance implies stationarity). If a Markov chain $\{\theta_j\}$ satisfies the detailed balance condition with π , then π is the stationary distribution.

4. (Convergence of Markov chain). Let $\{\theta_j\}$ be a Markov chain with $\theta_0 \sim \pi_{\text{initial}}(\cdot)$. If (i) π is the stationary distribution of the chain, and (ii) regularity conditions hold, then for large J

$$[\theta_J] \approx \pi(\cdot).$$

Remark 2.1. The above theorem suggests that, if we can design a suitable K so that the corresponding chain has a limiting distribution π , then when this chain is long enough, the samples from this chain can be regarded as samples from the target distribution.

Theorem 2.2. (Limiting behaviours of MCMC estimator)

1. (Ergodic theorem). Let g be a function such that $E_\pi |g(\theta)| < \infty$. Let $\{\theta_j\}$ be a Markov chain with $\theta_0 \sim \pi_{\text{initial}}(\cdot)$. Define

$$\hat{I}_{\text{MCMC}} := \frac{1}{J} \sum_{j=1}^J g(\theta_j) \quad \text{and} \quad I = E_\pi \{g(\theta)\}.$$

If (i) π is the stationary distribution of the chain, and (ii) regularity conditions hold, then as $J \rightarrow \infty$ $\hat{I}_{\text{MCMC}} \xrightarrow{\text{wp1}} I$.

2. (Central limit theorem). Let g be a function such that $E_\pi \{g^2(\theta)\} < \infty$. Let $\{\theta_j\}$ be a Markov chain with $\theta_0 \sim \pi_{\text{initial}}(\cdot)$. If π is the stationary distribution of $\{\theta_j\}$ and regularity conditions hold, then for any $\pi_{\text{initial}}(\cdot)$

$$\sqrt{J} \left(\hat{I}_{\text{MCMC}} - I \right) \xrightarrow{d} N(0, \sigma_{\text{MCMC}}^2)$$

where

$$\sigma_{\text{MCMC}}^2 = \sum_{k \in \mathbb{Z}} \gamma_k, \quad \text{and} \quad \gamma_k = \text{Cov}_\pi \{g(\theta_0), g(\theta_k)\}$$

2.2 Metropolis–Hastings (MH) Algorithm

Algorithm 2: Metropolis–Hastings Algorithm.

Input: (i) simulation size J ; (ii) proposed PDF $p(\cdot | \cdot)$; (iii) (unnormalized) target density $\pi_u(\cdot)$; (iv) initialization density $\pi_{\text{initial}}(\cdot)$.

begin

(1) Generate $\theta_0 \stackrel{\text{iid}}{\sim} \pi_{\text{initial}}(\cdot)$.

(2) **for** j in $\{1, \dots, J\}$ **do**

Generate $\tilde{\theta}_j \stackrel{\text{iid}}{\sim} p(\cdot | \theta_{j-1})$.

Generate $U_j \sim \text{Unif}(0, 1)$.

Compute the acceptance probability:

$$a_j = \min \left\{ 1, \frac{\pi_u(\tilde{\theta}_j)}{\pi_u(\theta_{j-1})} \frac{p(\theta_{j-1} | \tilde{\theta}_j)}{p(\tilde{\theta}_j | \theta_{j-1})} \right\}$$

$$= \min \left\{ 1, \text{Target odd} / \text{Proposal odd} \right\},$$

$$\text{Target odd} = \pi_u(\tilde{\theta}_j) / \pi_u(\theta_{j-1}),$$

$$\text{Proposal odd} = p(\tilde{\theta}_j | \theta_{j-1}) / p(\theta_{j-1} | \tilde{\theta}_j).$$

Compute the value in the j th iteration: $\theta_j = \tilde{\theta}_j \mathbb{1}(U_j \leq a_j) + \theta_{j-1} \mathbb{1}(U_j > a_j)$

end

end

Output: $\theta_{1:J}$

Remark 2.2. Some remark.

1. If the proposal is symmetric, i.e. $p(\tilde{\theta} | \theta) = p(\theta | \tilde{\theta})$, then the proposal odd can be dropped.
2. To avoid numerical issue, it is recommended to compute $a_j = \exp(a_j^*)$ where

$$a_j^* = \min\{0, \log \text{Target odd} - \log \text{Proposal odd}\},$$

$$\text{Target odd} = \log \pi_u(\tilde{\theta}_j) - \log \pi_u(\theta_{j-1}),$$

$$\text{Proposal odd} = \log p(\tilde{\theta}_j | \theta_{j-1}) - \log p(\theta_{j-1} | \tilde{\theta}_j).$$

3. Typically, we drop first half (depends on the situation) of the simulated samples and use the remaining for further analysis.

Theorem 2.3 (Justification of MH Algorithm). *The output from MH Algorithm satisfies the detailed balance condition with π .*

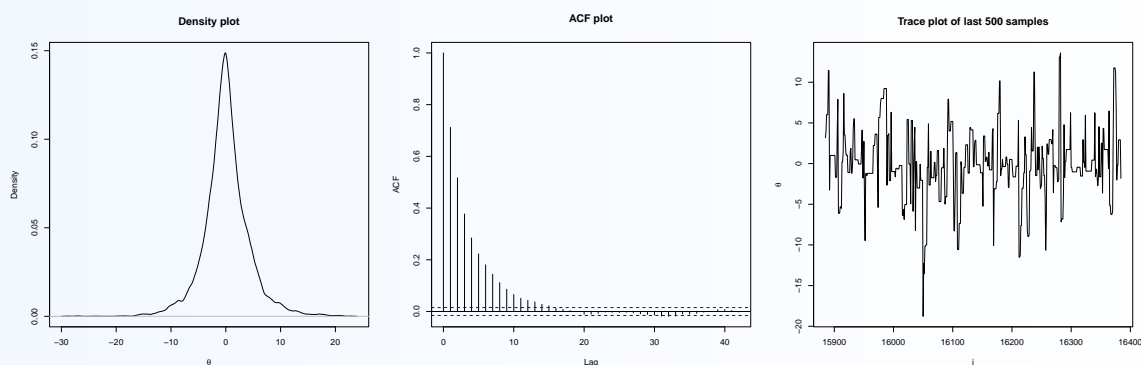
Example 2.1. Continue from Example 1.1, estimate $E\theta^2$ by the MH algorithm.

SOLUTION: Consider the initial proposal being $\text{Normal}(0, 8^2)$ and the proposal $\tilde{\theta} | \theta \sim \text{Normal}(\theta, 8^2)$. Note that the proposal is symmetric in this case.

```

1 ##MH algorithm
2 ##sd0 is the sd of the initial proposal
3 ##sd1 is the sd of the proposal
4 pi_u <- function(theta,b=3){
5   exp(-abs(theta)/b)
6 }
7
8 target_g <- function(theta){
9   theta^2
10 }
11
12 MH <- function(J,burn_frac = 0.5, keep = F,sd0 = 8, sd1 = 8,b = 3){
13   nsim = ceiling(J/burn_frac)+1
14   theta = array(NA,nsim)
15   theta[1] = rnorm(1,0,sd0) ##sample from the initial proposal
16   for (j in 2:nsim) {
17     tilde_theta = rnorm(1,theta[j-1],sd1) ##sample from the proposal
18     u = runif(1)
19     log_target_odd = log(pi_u(tilde_theta,b)) - log(pi_u(theta[j-1],b))
20     a = exp(min(0,log_target_odd))
21     theta[j] = tilde_theta*(u <= a) + theta[j-1]*(u > a)
22   }
23   if (keep) {
24     return(theta)
25   }else{
26     return(tail(theta,J)) ##burn the first fraction
27   }
28 }
29
30 J_sim = 2^14
31 theta_MH = MH(J_sim,sd0 = 8,sd1=8)
32
33 ##MCMC Diagostic
34 pi = density(theta_MH, kernel="epanechnikov")
35 acf = acf(theta_MH,plot=F)
36 par(mfrow = c(1,3))
37 plot(pi,main = 'Density plot',ylab = 'Density',xlab = expression(theta))
38 plot(acf,main = 'ACF plot',ylab = 'ACF',xlab = 'Lag')
39 plot((J_sim-500+1):J_sim,tail(theta_MH,500),type = 'l',
40      main = 'Trace plot of last 500 samples',ylab = expression(theta),xlab = '
41      j')
42 ##Produce estimate for I
43 mean(target_g(theta_MH))
44 [1] 17.82765

```



Example 2.2 (Option Pricing). Black-Scholes Formula is well known for the European option pricing. Specifically, we have the call price c being

$$\begin{aligned} c(\sigma) &= S_0 N(d_1) - K e^{-rT} N(d_2) \\ d_1 &= \frac{\log(S_0/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}} \\ d_2 &= d_1 - \sigma\sqrt{T}, \end{aligned}$$

where S_0 is current stock price, T is time to maturity, K is the strike price, σ is the volatility of the stock return R and $N(\cdot)$ is the CDF of the standard normal distribution. In practice, the σ is typically unknown. Consider following model,

$$\begin{aligned} R_{1:n} | \sigma &\stackrel{\text{iid}}{\sim} f(R | \theta) = \frac{1}{\sqrt{2}\sigma} \exp\left(-\frac{\sqrt{2}|R|}{\sigma}\right) \mathbb{1}(R \in \mathbb{R}) \\ \sigma &\sim Ga(\alpha)/\beta, \end{aligned}$$

where $\alpha = 0.4$ and $\beta = 2$.

1. Compute the posterior.
2. Using the MH algorithm, give a point estimate for σ .
3. Given the data with $n = 30$ and $A_n = \sum_{i=1}^n |R_i| = 2$. Suggest the call price when $S_0 = 373$, $K = 380$, $T = 0.5$ and $r = 0\%$.

SOLUTION:

1. By direct computation,

$$f(\sigma | R_{1:n}) \propto \sigma^{\alpha-n-1} \exp\left(-\beta\sigma - \frac{\sqrt{2}}{\sigma}A_n\right) \mathbb{1}(\sigma > 0),$$

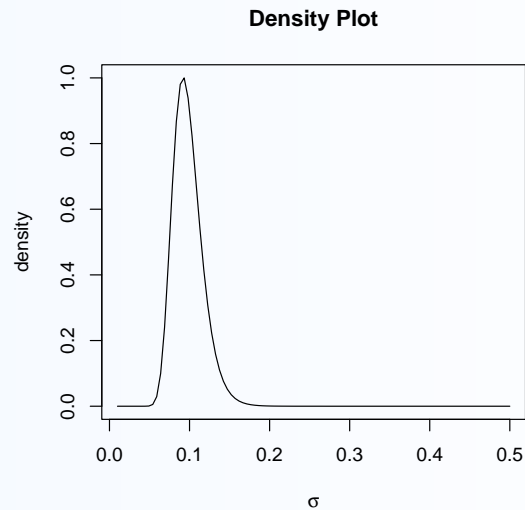
where $A_n = \sum_{i=1}^n |R_i|$.

2. We first visualize the target posterior density.

```

1 An = 2
2 n = 30
3
4 log_pi_u2 = function(sigma, An, n, a = 0.4, b = 2){
5   (a-n-1)*log(sigma) - b*sigma - sqrt(2)*An/sigma
6 }
7
8 ##visualize the target density
9 sigma_grid = seq(0.01, 0.5, length.out = 101)
10 log_den = log_pi_u2(sigma_grid, An, n)
11 plot(sigma_grid, exp(log_den - max(log_den)), type = 'l',
12       main = 'Density Plot', ylab = 'density', xlab = expression(sigma))

```



- The plot suggests that we may consider the Gamma distribution as initial point and the proposal.
- From the plot, it suggest that we should consider the initial proposal centers at $a/b = 0.1$ with standard deviation (sd) being $\sqrt{a/b} = 0.06$. Therefore, the initial proposal is $Ga(4)/40$.
- We want the proposal $[\tilde{\theta} | \theta]$ to center at θ while keeping the sd same as the initial. Therefore, the proposal is $Ga(400\theta^2)/400\theta$.

The MH algorithm can be implemented by the following code.

```

1 MH2 <- function(J,An,n,burn_frac = 0.5, keep = F,a0 = 4,b0=40){
2   nsim = ceiling(J/burn_frac)+1
3   theta = array(NA,nsim)
4   theta[1] = rgamma(1,a0,b0) ##sample from the initial proposal
5   for (j in 2:nsim) {
6     a1 = 400*theta[j-1]^2
7     b1 = 400*theta[j-1]
8     tilde_theta = rgamma(1,a1,b1) ##sample from the proposal
9     u = runif(1)
10    log_target_odd = log_pi_u2(tilde_theta,An,n) - log_pi_u2(theta[j-1],
11      An,n)
12    a2 = 400*tilde_theta^2
13    b2 = 400*tilde_theta
14    log_proposal_odd = log(dgamma(tilde_theta,a1,b1)) - log(dgamma(theta[
15      j-1],a2,b2))
16    a = exp(min(0,log_target_odd-log_proposal_odd))
17    theta[j] = tilde_theta*(u <= a) + theta[j-1]*(u > a)
18  }
19  if (keep) {
20    return(theta)
21  }else{
22    return(tail(theta,J)) ##burn the first fraction
23  }
24 }
J_sim = 2^14

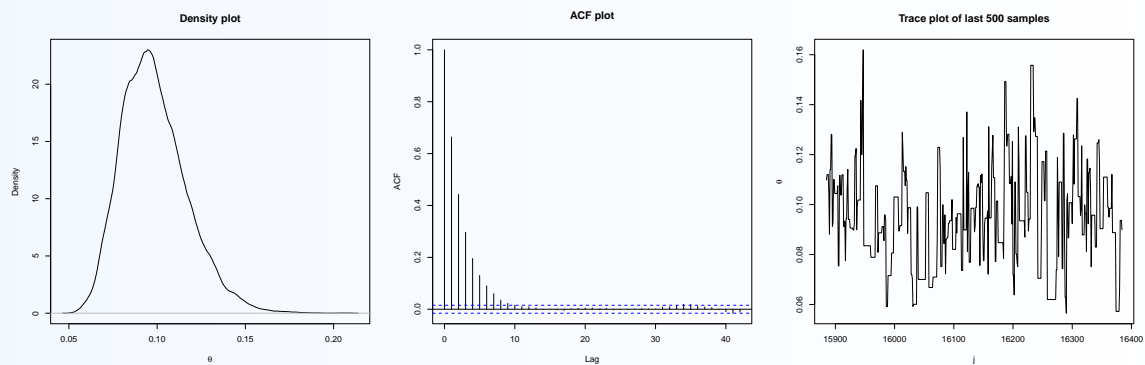
```



```

25 theta_MH2 = MH2(J_sim,An,n)
26
27 ##MCMC Diagostic
28 pi = density(theta_MH2, kernel="epanechnikov")
29 acf = acf(theta_MH2,plot=F)
30 windows(height = 5,width = 15)
31 par(mfrow = c(1,3))
32 plot(pi,main = 'Density plot',ylab = 'Density',xlab = expression(theta))
33 plot(acf,main = 'ACF plot',ylab = 'ACF',xlab = 'Lag')
34 plot((J_sim-500+1):J_sim,tail(theta_MH2,500),type = 'l',
35      main = 'Trace plot of last 500 samples',ylab = expression(theta),
36      xlab = 'j')
37 #alternative R functions to produce the same plot
38 par(mfrow = c(1,3))
39 plot(pi,main = 'Density plot',ylab = 'Density',xlab = expression(theta))
40 acf = acf(theta_MH2,plot=T)
41 ts.plot(tail(theta_MH2,500),type = 'l',
42        main = 'Trace plot of last 500 samples',ylab = expression(theta),
43        xlab = 'j')

```



3. Under the square loss, the Bayes estimator for the call price $c(\sigma)$ is posterior mean $E[c(\sigma) | R_{1:n}]$ estimated as \$7.2880.

```

1 ##Option pricing
2 get_call <- function(sigma_hat,s0 = 373,k = 380,t = 0.5,r = 0){
3   d1 = (log(s0/k)+(r+sigma_hat^2/2)*t)/(sigma_hat*sqrt(t))
4   d2 = d1-sigma_hat*sqrt(t)
5   c = s0*pnorm(d1)-k*exp(-r*t)*pnorm(d2)
6   c
7 }
8 mean(get_call(theta_MH2))
9 [1] 7.287991

```

Remark 2.3. For a comprehensive procedure of using different methods on the same task (together with R codes), you may study Example 7.1 (the model), Example 7.2 (Trapezoidal rule), Example 7.4 (Importance Sampling) 7.6 (Reweighting Importance Sampling) and Example 7.7 (MH Algorithm) in details. It will help assignment 6. However, other examples from the lecture notes are also important.