

## Lab 12 of Thursday 4th December 2025

### Exercise 1.

At every iteration of the general Metropolis–Hastings algorithm, a new candidate state  $Y_{n+1}$  is proposed by sampling  $Y_{n+1} \sim q(X_n, \cdot)$ , given the current state  $X_n$ . Here,  $q(x, y)$  is the so-called proposal density. Consider now the case where the proposal does not depend on the current state, that is  $q(x, y) \equiv q(y)$ , so that the proposed candidate is  $Y_{n+1} \sim q$ . This particular Markov Chain Monte Carlo (MCMC) variant is sometimes called *independent Metropolis–Hastings algorithm* with fixed proposal (or simply *independence sampler*). Let's denote the target density by  $f$ . As such, this MCMC variant appears very similar to the Accept–Reject method for sampling from  $f$  (cf. Lab 02).

- 1) Suppose there exists a positive constant  $C$  such that  $f(x) \leq Cq(x)$  for any  $x \in \text{supp}(f) = \{x \in \mathbb{R}^d : f(x) > 0\}$ . Show that the expected acceptance probability of the independent Metropolis–Hastings algorithm is *at least*  $\frac{1}{C}$  whenever the chain is stationary. How does this compare to the expected acceptance probability of an Accept–Reject method?
- 2) Let us compare the independent Metropolis–Hastings algorithm and the Accept–Reject method in some more detail by an example. Specifically, the goal is to sample from a Gamma distribution with shape parameter  $\alpha$  and scale parameter  $\beta$ , denoted by  $\text{Gamma}(\alpha, \beta)$ , so that the target PDF reads  $f(x) \equiv f(x; \alpha, \beta) = \beta^\alpha x^{\alpha-1} e^{-\beta x} / \Gamma(\alpha) \mathbb{I}_{\{x \geq 0\}}$ , where  $\Gamma(\cdot)$  denotes the Gamma function.
  - a) Implement the Accept–Reject method to sample from  $\text{Gamma}(\alpha, 1)$  for  $\alpha > 1$ , using the PDF of the  $\text{Gamma}(a, b)$  distribution with  $a = [\alpha]$  as auxiliary density (here  $[\alpha]$  denotes the integer part of  $\alpha$ ).<sup>1</sup> Show that  $b = [\alpha]/\alpha$  is the optimal choice for  $b$ .
  - b) Use your Accept–Reject method to generate  $m$  random numbers  $X_1, \dots, X_m$  with each  $X_i \sim \text{Gamma}(\alpha, 1)$ , when using  $n = 5000$  random variables  $Y_1, \dots, Y_n$  from the auxiliary  $\text{Gamma}([\alpha], [\alpha]/\alpha)$  distribution. Notice that  $m$  is a random variable, which is smaller than  $n$  due to rejections. Perform the simulations for  $\alpha = 4.85$ .
  - c) Implement the independent Metropolis–Hastings algorithm using as proposal  $q$  the PDF of the  $\text{Gamma}([\alpha], [\alpha]/\alpha)$  distribution.
  - d) Use the same sample  $Y_1, \dots, Y_n$  used within the Accept–Reject method, now in the corresponding Metropolis–Hastings algorithm to generate  $n = 5000$  realizations of the target distribution  $\text{Gamma}(\alpha, 1)$  with  $\alpha = 4.85$ .
  - e) Compare both methods with respect to:
    - i. their acceptance rates,
    - ii. their estimates for the mean of the  $\text{Gamma}(4.85, 1)$  distribution, which is 4.85,

---

<sup>1</sup>Hint: Recall that  $\sum_{k=1}^K \xi_k \sim \text{Gamma}(K, \beta)$  for  $K \in \mathbb{N}$ , if  $\xi_k \stackrel{\text{i.i.d.}}{\sim} \text{Gamma}(1, \beta) \equiv \text{Exp}(\beta)$ .

- iii. the correctness of the target distribution,  
 Discuss your results.

### Solution

- 1) Let  $C > 0$  such that  $f(x) \leq Cq(x)$  for any  $x \in \text{supp}(f)$ . Suppose that the chain is stationary. Then the expected acceptance probability is

$$\begin{aligned} \mathbb{E} \left( \min \left\{ \frac{f(Y_{n+1})q(X_n)}{f(X_n)q(Y_{n+1})}, 1 \right\} \right) &= \int \mathbb{I}_{\left\{ \frac{f(y)q(x)}{q(y)f(x)} \geq 1 \right\}} f(x)q(y) \, dx \, dy \\ &+ \int \frac{f(y)q(x)}{q(y)f(x)} \mathbb{I}_{\left\{ \frac{f(y)q(x)}{q(y)f(x)} < 1 \right\}} f(x)q(y) \, dx \, dy \\ &= 2 \int \mathbb{I}_{\left\{ \frac{f(y)q(x)}{q(y)f(x)} \geq 1 \right\}} f(x)q(y) \, dx \, dy \\ &\geq 2 \int \mathbb{I}_{\left\{ \frac{f(y)}{q(y)} \geq \frac{f(x)}{q(x)} \right\}} f(x) \frac{f(y)}{C} \, dx \, dy \\ &= \frac{2}{C} \mathbb{P} \left( \frac{f(X_1)}{q(X_1)} \geq \frac{f(X_2)}{q(X_2)} \right) = \frac{1}{C}, \end{aligned}$$

where the last equality follows from the fact that  $X_1$  and  $X_2$  are independent and both distributed according to  $f$  by assumption. In contrast, the average acceptance probability for an Accept–Reject method is always equal to  $1/C$ .

- 2) Proving that  $b = [\alpha]/\alpha$  is the optimal choice for  $b$  follows from straightforward calculations. In fact, the ratio  $f(x; \alpha, 1)/f(x; a, b)$  is

$$\frac{f(x; \alpha, 1)}{f(x; a, b)} = b^{-a} x^{\alpha-a} e^{-(1-b)x} \frac{\Gamma(a)}{\Gamma(\alpha)},$$

for any  $x \geq 0$ , which yields the bound  $C = b^{-a} \left( \frac{\alpha-a}{(1-b)e} \right)^{\alpha-a}$  for  $b < 1$ . The optimal choice for  $b$  then follows by optimization. A possible `Python` implementation for this exercise is shown at the end of this Section. The code produced, for example the following output regarding the acceptance rate:

```
-----
AR acceptance rate 0.9053050878145935
MH acceptance rate 0.94
-----

AR mean 4.835818571907215
MH mean 4.820798409910968
-----
```

confirming the theoretical result proved above. Furthermore, the code also produced the plots shown in Figures 1 and 2.

These plots show that both methods produce accurate approximations to the mean and the PDF. Moreover, we would expect to see that the sample generated by the MH algorithm




Figure 1: (Left) Kernel density estimate (KDE) Metropolis Hastings. (Middle) KDE accept-reject method. Both shown on top of true density. (Right) Ergodic estimator of the mean for both methods.




Figure 2: Autocorrelation plot of AR (left) and MH (right). The blue band represents a 95% confidence interval.

contains more correlations, this is because the the AR method should, in fact, produce independent samples, while the MH is a Markov chain. This, however, is difficult to appreciate in the autocorrelation plots, as the samples obtained using MH de-correlate quite rapidly.

### Python code:

```
import numpy as np
import scipy.special as sps
import matplotlib.pyplot as plt
from pandas import Series
import statsmodels.graphics.tsaplots as sm
import seaborn as sns
sns.set(color_codes=True)# Defines the pdf
sns.set(font_scale=2) #fontsize in plots
sns.set_style("white")

np.random.seed(42)

def f(x,a,b):
    return b**a*x**(a-1)*np.exp(-b*x)/sps.gamma(a)*1*(x>=0)

def Cfunc(alpha, a, b):
    return ((alpha - a)/(1-b))**(alpha-a) * np.exp(a-alpha) * sps.gamma(a)/sps.gamma(alpha)/b**a

#### Some parameters
alpha=4.85
a = np.floor(alpha)
b = np.floor(alpha)/alpha
C = Cfunc(alpha,a,b)

#### Does the accept reject part
n = 5000
Aar = 0
Xar = np.zeros(n)
i = 0
p = 0
while i<n:
    xi = np.random.gamma(shape = a, scale = 1/b, size=1)
    ratio = f(xi,alpha,1)/(f(xi,a,b)*C)
    p+=1
    if np.random.random(1) < ratio:
        Xar[i]=xi
        i+=1
Aar=n/p # acceptance probability

### Does the Metropolis Hastings part
Xmh = np.zeros(n)
Xmh[0] = 4
Amh=0
for i in range(n-1):
    #samples proposal
    y = np.random.gamma(shape = a, scale = 1/b, size=1)
    #computes MH ratio
    py=f(y,alpha,1)
    px=f(Xmh[i],alpha,1)
    qyx=f(Xmh[i],a,b)
    qxy=f(y,a,b)

    ratio=(py*qyx)/(px*qxy)
    if np.random.random(1) < ratio:
        Xmh[i+1]=y
        Amh+=1
```

```

else:
    Xmh[i+1]=Xmh[i]
Amh=Amh/n

print('-----')
print('AR acceptance rate '+str(Aar))
print('MH acceptance rate '+str(Amh))
print('-----')
print('AR mean '+str(np.mean(Xar)))
print('MH mean '+str(np.mean(Xmh)))
print('-----')

xx=np.linspace(0,20,1000)
plt.plot(xx,C*f(xx,a,b))
plt.plot(xx,f(xx,alpha,1))
plt.legend(["Cg(x)", "f(x)"] )
plt.savefig("../figures/densities.png")
plt.show()

nn=np.arange(1,n+1)
plt.plot(nn,np.cumsum(Xar)/nn)
plt.plot(nn,np.cumsum(Xmh)/nn)
plt.hlines(alpha,0,n,color='black')
plt.legend(["AR", "MH", "True mean"] )
plt.show()

sm.plot_acf(Xar,lags=20)
plt.gca().set_rasterized(True)
plt.legend(["AR"] )
plt.savefig("../figures/acfar.png")
plt.show()

sm.plot_acf(Xmh,lags=20)
plt.gca().set_rasterized(True)
plt.legend(["MH"] )
plt.show()

sns.kdeplot(Xar, shade=True, color="r")
plt.gca().set_rasterized(True)
x = np.linspace(0,20,100)
plt.plot(x, f(x,alpha,1),color='black')
plt.title('KDE for AR')
plt.show()

sns.kdeplot(Xmh, shade=True, color="r")
plt.gca().set_rasterized(True)
plt.plot(x, f(x,alpha,1),color='black')
plt.title('KDE for MH')
plt.show()

```

## Exercise 2.

Consider the Random Walk Metropolis–Hastings (RWMH) algorithm with proposal density  $q(x, y) = g_\sigma(y - x)$  and target density  $f: \mathbb{R} \rightarrow \mathbb{R}^+$ . Let  $g_\sigma$  denote the density of the  $\mathcal{N}(0, \sigma^2)$  distribution and suppose that

$$f(y) = \frac{1}{Z} \exp \left[ - \left( \frac{1}{4} y^4 - \frac{1}{2} y^2 + \frac{1}{4} \right) \right],$$

where  $Z$  is such that  $f$  is a PDF on  $\mathcal{X} = \mathbb{R}$ . Suppose we wish to estimate  $\mu = \mathbb{E}_f(\phi)$  for a suitable function  $\phi: \mathbb{R} \rightarrow \mathbb{R}$ . Let  $\hat{\mu}_n^{\text{MH}}$  be the estimator for  $\mu$  based on the Markov chain of length  $n$  generated by the RWMH algorithm. Derive asymptotic confidence intervals for  $\mu$  at probability level  $\alpha$  using the CLT for Metropolis–Hastings Markov Chains. Within your simulations, estimate<sup>2</sup> this confidence interval and stop the Markov chain once the half-length of the interval is smaller than a given tolerance  $\tau > 0$ . Implement the following heuristics to estimate the required *time average variance constant* and compare their performance for different functions  $\phi$ ; namely  $\phi(x) = x^p, p \in \mathbb{N}$ . The time average variance constant is the asymptotic variance introduced in Theorem 8.10 of the lecture notes.

1) *Initial positive sequence estimator:*

$$\tilde{\sigma}^2 \approx \hat{\sigma}_{\text{pos},n}^2 := -\hat{c}_n(0) + 2 \sum_{k=1}^K (\hat{c}_n(2k) + \hat{c}_n(2k+1)),$$

where  $K$  is the largest integer such that  $\hat{c}_n(2k) + \hat{c}_n(2k+1) > 0$  for all  $k = 1, \dots, K$ . Here,

$$\hat{c}_n(j) := \frac{1}{n} \sum_{i=1}^{n-j} (\phi(X_i) - \hat{\mu}_n^{\text{MH}}) (\phi(X_{i+j}) - \hat{\mu}_n^{\text{MH}})$$

is an appropriate covariance estimator in this context.

2) *Initial monotone sequence estimator:*

$$\tilde{\sigma}^2 \approx \hat{\sigma}_{\text{mon},n}^2 := -\hat{c}_n(0) + 2 \sum_{k=1}^K \min_{1 \leq j \leq k} \{\hat{c}_n(2j) + \hat{c}_n(2j+1)\},$$

where  $K$  and  $\hat{c}_n$  are as for the initial positive sequence estimator above.

3) *Batch means estimator:* Suppose the Markov chain is  $X_1, \dots, X_n$  at iteration  $n$ . Divide these  $n$  values into  $N_b \in \mathbb{N}$  batches, each of length  $N_\ell = n/N_b$ . A typical decomposition is  $N_\ell = n^{1-a}$  and  $N_b = n^a$  for  $a \in [0, 1]$ , for example  $a = 0.5$ , modulo integer rounding. Let

$$\hat{\mu}_i = \frac{1}{N_\ell} \sum_{j=(i-1)N_\ell+1}^{iN_\ell} \phi(X_j), \quad i = 1, \dots, N_b,$$

be the sample mean of the  $i$ -th batch. For  $N_\ell$  sufficiently large, one can consider the batch means  $\hat{\mu}_1, \hat{\mu}_2, \dots, \hat{\mu}_{N_b}$  to be approximately mutually independent. Consequently, one can estimate the time average variance constant  $\sigma^2$  by the sample variance estimator for independent realizations:

$$\tilde{\sigma}^2 \approx \hat{\sigma}_{\text{BM},n}^2 := \frac{n}{N_b} \frac{1}{N_b - 1} \sum_{i=1}^{N_b} (\hat{\mu}_i - \hat{\mu}_n^{\text{MH}})^2.$$

In addition, instead of using all of the points in the Markov chain, one can as well discard a burn-in time  $B$  and replace  $\sum_{k=1}$  with  $\sum_{k=B}$  in the above estimators. Experiment with the effects of burn-in time on the above asymptotic variance estimators.

## Solution

<sup>2</sup>The estimation has to be carried out on-the-fly, that is while the Markov chain evolves.

## Solution

An asymptotic confidence interval can be derived in a straightforward manner from Theorem 8.17 of the lecture notes. Let  $\{X_n\}_{n=1}^N$  denote the generated Markov chain and  $\hat{\mu}_N^{mcmc}$  denote the standard MCMC estimator of  $\mu = \mathbb{E}[\phi]$ . We know that asymptotically as  $N \rightarrow \infty$ ,

$$\sqrt{N}(\hat{\mu}_N^{mcmc} - \mu) \xrightarrow{d} \mathcal{N}(0, \sigma_{mcmc}^2), \quad (2.1)$$

where  $\sigma_{mcmc}^2$  denotes the asymptotic variance or time average variance constant of the Markov chain. A  $1 - \alpha$  confidence interval then simply reads  $I = \left[ \mu \pm \frac{C_{1-\alpha/2} \sigma_{mcmc}}{\sqrt{N}} \right]$  where  $C_\delta$  denotes the inverse of the CDF of the standard normal distribution evaluated at  $\delta$ .

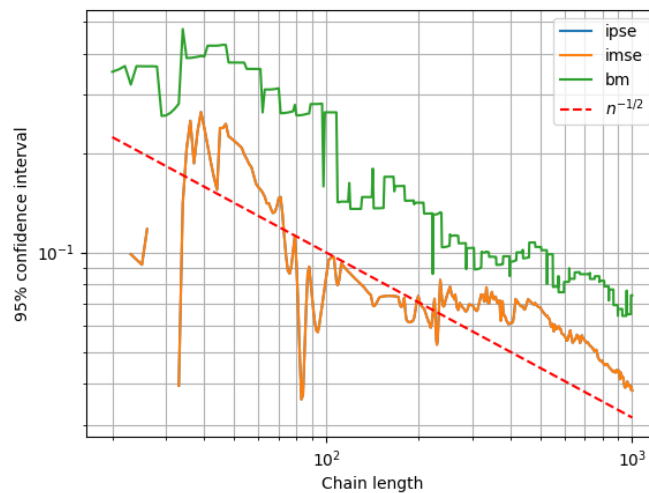


Figure 3: Convergence for different asymptotic variance estimators

### Python code:

```
import matplotlib.pyplot as plt
import numpy as np
import scipy.stats as st
import statsmodels.api as sm

np.random.seed(42)

def f(y):
    return np.exp(-((y**2 - 0.5) ** 2))

def get_K(auto_covars):
    K = None
    for k in range(int(len(auto_covars) / 2)):
        if (auto_covars[2 * k] + auto_covars[2 * k + 1]) <= 0:
            K = k
            break
    if K is None:
        K = int(len(auto_covars) / 2 - 1)
    return K
```

```

def sigma_ipse(auto_covars):
    K = get_K(auto_covars)
    return -auto_covars[0] + 2 * np.sum(
        [auto_covars[2 * k] + auto_covars[2 * k + 1] for k in range(1, K + 1)]
    )

def sigma_imse(auto_covars):
    K = get_K(auto_covars)
    return -auto_covars[0] + 2 * np.sum(
        [
            np.min(
                [auto_covars[2 * j] + auto_covars[2 * j + 1] for j in range(1, k + 1)]
            )
            for k in range(1, K + 1)
        ]
    )

def sigma_bm(Xchain, a=0.5):
    n = len(Xchain)
    Nb = int(np.floor(n**a))
    Nl = int(np.floor(n / Nb))
    batch_means = []
    mcmc_mean = np.mean(Xchain)
    for i in range(1, Nb + 1):
        batch_means.append(np.mean([Xchain[j] for j in range((i - 1) * Nl, i * Nl)]))

    batch_means = np.array(batch_means)
    return np.sum((batch_means - mcmc_mean) ** 2) * n / (Nb - 1) / Nb

def compute_auto_covars(Xchain):
    Xmean = np.mean(Xchain)
    N = len(Xchain)
    autocov = np.zeros(N)
    for k in range(N):
        for i in range(N - k):
            autocov[k] += ((Xchain[i + k]) - Xmean) * (Xchain[i] - Xmean)
        autocov[k] = (1 / (N - 1)) * autocov[k]
    return autocov

# Simulation parameters
alpha = 0.05
coeff = st.norm.ppf(1 - alpha / 2)
sigma = 4.0
p = 2

# Start the chain
Xchain = np.array([0])
n = 0
error1s = []
error2s = []
error3s = []

iteration = 0
iter_min = 20 # force a minimum number of iteration before checking stopping criterion
iter_stop = 1000 # max iteration limit

ns = []

```

```

while True:
    # Compute MH step
    Xn = Xchain[n]
    Xnew = st.norm.rvs() * sigma
    uniform = st.uniform.rvs()
    ratio = np.min([1, (f(Xnew) * st.norm.pdf(Xn)) / (f(Xn) * st.norm.pdf(Xnew))])
    if uniform <= ratio:
        Xchain = np.append(Xchain, Xnew)
    else:
        Xchain = np.append(Xchain, Xn)
    n = n + 1

if n >= iter_min:
    # Compute phi(X)
    Xchain_p = Xchain**p

    # Compute autocovariances
    # Autocovariances computed from scratch, not incrementally.
    # Can be reimplemented for speed
    auto_covars = compute_auto_covars(Xchain_p)

    # Compute variances
    sigma1 = sigma_ipse(auto_covars)
    sigma2 = sigma_imse(auto_covars)
    sigma3 = sigma_bm(Xchain_p)

    error1 = coeff * np.sqrt(sigma1) / np.sqrt(n)
    error2 = coeff * np.sqrt(sigma2) / np.sqrt(n)
    error3 = coeff * np.sqrt(sigma3) / np.sqrt(n)

    print(n, error1, error2, error3)
    error1s.append(error1)
    error2s.append(error2)
    error3s.append(error3)

    ns.append(n)
if n >= iter_stop:
    break

plt.loglog(ns, error1s, label="ipse")
plt.loglog(ns, error2s, label="imse")
plt.loglog(ns, error3s, label="bm")
plt.loglog(ns, np.array(ns) ** -0.5, "--", color="red", label=r"$n^{-1/2}$")
plt.grid(which="both")
plt.xlabel("Chain length")
plt.ylabel("95% confidence interval")
plt.legend()
plt.show()

```

### Exercice 3.

Consider a Markov chain  $\{X_n\} \sim \text{Markov}(\pi, P)$  on a discrete state space  $\mathcal{X}$  at equilibrium, with  $P$  irreducible, and  $\pi$  the unique invariant probability measure of  $P$ . Let  $l_\pi^2$  be the Hilbert space  $l_\pi^2 = \{f : \mathcal{X} \rightarrow \mathbb{R} : \sum_{i \in \mathcal{X}} f(i)^2 \pi_i < \infty\}$  with inner product  $(f, g)_{l_\pi^2} = \sum_{i \in \mathcal{X}} f(i)g(i)\pi_i$ , and  $l_{\pi,0}^2 = \{f \in l_\pi^2 : \mathbb{E}_\pi[f] = 0\}$ .

- 1) Show that if  $(P, \pi)$  are in detailed balance, then  $(Pf, g)_{l_\pi^2} = (f, Pg)_{l_\pi^2}$  for any  $f, g \in l_\pi^2$
- 2) Show that  $\mathbb{E}[f(X_n)f(X_m)] = (P^{m-n}f, f)_{l_\pi^2}$  for any  $f \in l_\pi^2$  and  $m > n$ .

3) Consider now the estimator

$$\hat{\mu}_N = \frac{1}{N} \sum_{n=1}^N f(X_n)$$

of  $\mu = \mathbb{E}_\pi[f]$  under the assumption that  $f \in l_\pi^2$ . Show that  $\mathbb{E}_\pi[\hat{\mu}_N] = \mu$ , and

$$\text{Var}[\hat{\mu}_N] = \frac{1}{N} \sum_{l=0}^N c_l (P^l \tilde{f}, \tilde{f})_{l_\pi^2},$$

with  $\tilde{f} = f - \mathbb{E}_\pi[f] \in l_{\pi,0}^2$  and

$$c_{l,N} = \begin{cases} 1, & l = 0 \\ 2(1 - \frac{l}{N}), & l > 0 \end{cases} \quad (3.1)$$

4) Conclude that the asymptotic variance  $\mathbb{V}(f, p) := \lim_{N \rightarrow \infty} N \text{Var}_\pi(\hat{\mu}_N)$  satisfies  $\mathbb{V}(f, p) = ((2(I - P)^{-1} - I)\tilde{f}, \tilde{f})_{l_\pi^2}$  if

$$\sup_{g \in l_{\pi,0}^2} \frac{(Pg, g)_{l_\pi^2}}{\|g\|_{l_\pi^2}} = \beta < 1. \quad (3.2)$$

5) Consider now the two irreducible transition matrices  $P_1$  and  $P_2$ , both in detailed balance with  $\pi$  and satisfying (3.2) for some  $\beta_1, \beta_2$ . Show that if  $(P_1)_{ij} \geq (P_2)_{ij} \forall i \neq j$ , then

$$\mathbb{V}(f, P_1) \leq \mathbb{V}(f, P_2), \quad (3.3)$$

for any  $f \in l_\pi^2$ .

**Hint:** Take  $P(\lambda) = (1 - \lambda)P_1 + \lambda P_2, \lambda \in [0, 1]$  and show that  $\frac{d}{d\lambda} \mathbb{V}(f, P(\lambda)) \geq 0$ .

## Solution