

# MONTE CARLO AND MARKOV CHAIN

## MONTE CARLO METHODS

**History:** Monte Carlo (MC) and Markov Chain Monte Carlo (MCMC) have been around for a long time. Some (very) early uses of MC ideas:

- Conte de Buffon (1777) dropped a needle of length  $L$  onto a grid of parallel lines spaced  $d > L$  apart to estimate  $P[\text{needle intersects a line}]$ .
- Laplace (1786) used Buffon's needle to evaluate  $\pi$ .
- Gosset (Student, 1908) used random sampling to determine the distribution of the sample correlation coefficient.
- von Neumann, Fermi, Ulam, Metropolis (1940s) used games of chance (hence MC) to study models of atomic collisions at Los Alamos during WW II.

We are concerned with the use of MC and, in particular, MCMC methods to solve estimation and detection problems.

As discussed in the introduction to MC methods (handout # 4), many estimation and detection problems require evaluation of integrals.

# Monte Carlo Integration

MC Integration is essentially numerical integration and thus may be thought of as estimation — we will discuss MC estimators that estimate integrals.

Although MC integration is most useful when dealing with high-dimensional problems, the basic ideas are easiest to grasp by looking at 1-D problems first.

Suppose we wish to evaluate

$$G = \int_{\Omega} g(x) p(x) dx$$

where  $p(\cdot)$  is a density, i.e.  $p(x) \geq 0$  for  $x \in \Omega$  and  $\int_{\Omega} p(x) dx = 1$ . Any integral can be written in this form if we can transform their limits of integration to  $\Omega = (0, 1)$  — then choose  $p(\cdot)$  to be uniform(0, 1).

The basic MC estimate of  $G$  is obtained as follows:

1. Draw i.i.d. samples  $x_1, x_2, \dots, x_N$  from  $p(\cdot)$  and
2. Estimate  $G$  as

$$\hat{G}_N = \frac{1}{N} \sum_{i=1}^N g(x_i).$$

Clearly,  $\hat{G}_N$  will be such that

$$\mathbb{E}[\hat{G}_N] = G \quad \text{and} \quad \hat{G}_N \xrightarrow{p} G$$

which follows by applying the law of large numbers (LLN).

### Comments:

- When we get to MCMC, we will do essentially the same thing except that independence will not hold for the samples  $x_1, x_2, \dots, x_N$ . We then need to rely on ergodic theorems rather than LLN.
- $\hat{G}_N$  has rate of convergence  $N^{-1/2}$  which is
  - (i) slow (quadrature may have  $\sim N^{-4}$  convergence)
  - (ii) more efficient than quadrature or finite-difference methods when the dimensionality of integrals is large ( $> 6-8$ )

Define

$$\sigma^2 = \underbrace{\int g^2(x) p(x) dx - \left[ \int g(x) p(x) dx \right]^2}_{\text{var}(G(X_i))}.$$

The overall efficiency of MC calculation is proportional to  $t \sigma^2$  where  $t$  is the time required to sample  $x$  from  $p(x)$ .

# Bias and Variance in MC Estimators

Assume

$$\hat{G}_N = \frac{1}{N} \sum_{i=1}^N g(x_i).$$

for

$$G = \int_{\Omega} g(x) p(x) dx$$

where  $x_i \sim$  i.i.d. with pdf  $p(x)$ . As before

$$\mathbb{E}[\hat{G}_N] = G \quad \text{and} \quad \hat{G}_N \xrightarrow{P} G.$$

Also

$$\mathbb{E}[(\hat{G}_N - G)^2] = \mathbb{E}[(\hat{G}_N - \mathbb{E}[\hat{G}_N])^2] + (\mathbb{E}[\hat{G}_N] - G)^2$$

which is just  $\text{MSE} = \text{variance} + \text{bias}^2$  (recall handout # 1). As we know from the estimation theory, it might be possible to find an estimator with smaller MSE than  $\hat{G}_N$  at the expense of being biased.

**Example.** Estimate the mean of  $\text{uniform}(0, 1)$  using MC:

$$G = \int_0^1 x dx, \quad \hat{G}_N^{(1)} = \frac{1}{N} \sum_i x_i$$

for  $x_i$  sampled from  $\text{uniform}(0, 1)$ . Consider

$$\hat{G}_N^{(2)} = \frac{1}{2} \max\{x_1, x_2, \dots, x_N\}$$

for  $x_i$  sampled from  $\text{uniform}(0, 1)$ .

$$\mathbb{E}[\hat{G}_N^{(1)}] = G \quad \text{whereas} \quad \mathbb{E}[\hat{G}_N^{(2)}] = \frac{N}{N+1} G.$$

But

$$\text{MSE}(\hat{G}_N^{(1)}) = \text{var}(\hat{G}_N^{(1)}) = \frac{G}{6N} \quad (\text{order } N^{-1})$$

$$\text{MSE}(\hat{G}_N^{(2)}) = \frac{2G^2}{(N+1)(N+2)} \quad (\text{order } N^{-2}).$$

### Comment:

- Of course, we may be able to get a better (and more complicated) unbiased estimator of  $G$  than  $\hat{G}_N^{(1)}$ . Yet, the biased estimator  $\hat{G}_N^{(2)}$  that performs well is really simple.

We often arrange for MC estimators to be unbiased, but this is not a necessary consequence of MC methodology.

**Example.** Formulate an MC estimator which samples from

uniform(0, 1) for the following integral:

$$G = \frac{\int_0^1 g_1(x) dx}{\int_0^1 g_2(x) dx}.$$

An estimate might be

$$\hat{G}_N = \frac{\sum_{i=1}^N g_1(x_i)}{\sum_{i=1}^N g_2(x_i)}$$

where  $x_1, x_2, \dots, x_N \sim$  i.i.d. uniform(0, 1). Here,  $\hat{G}_N$  will be biased for  $G$  although it will be consistent with  $\hat{G}_N \xrightarrow{P} G$ .

If  $\hat{G}_N$  is *unbiased*, we wish to reduce its variance (which is equal to the MSE in this case).

To implement MC methods we

1. must be able to generate from  $p(x)$  (or from a suitable alternative),
2. must be able to do so in reasonable time (i.e. small  $t$ ),
3. need small  $\sigma^2$ .

Monte Carlo addresses problem 3.

Markov Chain addresses problems 1 and 2.

# MC Variance Reduction: Importance Sampling

We have an integral

$$G = \int g(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}$$

where  $\mathbf{x}$  is  $p$ -dimensional.  $p(\mathbf{x})$  may not be the best distribution to sample from. Certainly,  $p(\mathbf{x})$  is not the only distribution we could sample from: For any  $\tilde{p}(\mathbf{x}) > 0$  *over the same support* as  $p(\mathbf{x})$ , satisfying  $\int \tilde{p}(\mathbf{x}) d\mathbf{x} = 1$ :

$$G = \int g(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} = \int \frac{g(\mathbf{x}) p(\mathbf{x})}{\tilde{p}(\mathbf{x})} \tilde{p}(\mathbf{x}) d\mathbf{x}$$

as long as

$$\frac{g(\mathbf{x}) p(\mathbf{x})}{\tilde{p}(\mathbf{x})} < \infty$$

up to a countable set and the original integral exists. Let  $\hat{G}_{\tilde{p}, N}$  be the MC estimator of  $G$  based on sampling from  $\tilde{p}$ :

$$\hat{G}_{\tilde{p}, N} = \frac{1}{N} \sum_i \frac{g(\mathbf{x}_i) p(\mathbf{x}_i)}{\tilde{p}(\mathbf{x}_i)} \quad (1)$$

where  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N \sim$  i.i.d.  $\tilde{p}(\mathbf{x})$ . Clearly,

$$\mathbb{E} [\hat{G}_{\tilde{p}, N}] = \mathbb{E} [\hat{G}_N] = G$$

for any legitimate  $\tilde{p}$  and

$$\text{var}(\hat{G}_{\tilde{p},N}) = \frac{1}{N} \int \left[ \frac{g(\mathbf{x})p(\mathbf{x})}{\tilde{p}(\mathbf{x})} \right]^2 \tilde{p}(\mathbf{x}) d\mathbf{x} - \frac{G^2}{N}.$$

The idea is to pick  $\tilde{p}$  to minimize this variance, subject to the constraint that  $\tilde{p}$  is a density.

Suppose that we pick

$$\tilde{p}(\mathbf{x}) \propto |g(\mathbf{x})| p(\mathbf{x})$$

i.e.

$$\tilde{p}(\mathbf{x}) = \frac{1}{C} |g(\mathbf{x})| p(\mathbf{x}) \quad \text{where} \quad C = \int |g(\mathbf{x})| p(\mathbf{x}) d\mathbf{x}.$$

This would give

$$\text{var}(\hat{G}_{\tilde{p},N}) = \frac{1}{N} \int C |g(\mathbf{x})| p(\mathbf{x}) d\mathbf{x} - \frac{G^2}{N} = \frac{1}{N} (C^2 - G^2).$$

Note that

if  $g(\mathbf{x}) > 0$  for all  $\mathbf{x}$  in the support of  $p$ , then

$$C = \int g(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} = G$$



and if  $g(\mathbf{x}) < 0$  for all  $\mathbf{x}$  in the support of  $p$ , then

$$C = - \int g(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} = -G.$$

In both cases  $\text{var}(\widehat{G}_{\tilde{p},N}) = 0!$

However, if we could compute  $\int |g(\mathbf{x})| p(\mathbf{x}) d\mathbf{x}$ , we would already be done and would not use MC. Nevertheless, we can draw some conclusions from the above exercise:

1. There may exist an importance pdf  $\tilde{p}(\mathbf{x})$  that gives smaller variance than  $p(\mathbf{x})$  (of the corresponding estimate of  $G$  in (1)).
2. Good  $\tilde{p}(\mathbf{x})$  are those that match the behavior of  $g(\mathbf{x}) p(\mathbf{x})$ . This is particularly true near the maximum of the integrand  $g(\mathbf{x}) p(\mathbf{x})$ .

The distribution  $\tilde{p}(\mathbf{x})$  we actually sample from is usually called the *importance function* or importance density.

How can we *actually* find useful importance densities?

**Example 1:** Find an MC estimator of

$$G = \int_0^1 (1 - x^2)^{1/2} dx.$$

Obviously, we could draw i.i.d. samples  $x_1, x_2, \dots, x_N$  from  $\text{uniform}(0, 1)$  and use

$$\hat{G}_N = \frac{1}{N} \sum_{i=1}^N (1 - x_i^2)^{1/2}.$$

It turns out that in this case

$$\text{var}(\hat{G}_N) = \frac{0.050}{N}.$$

Now, consider finding a good importance density. Observe that  $(1 - x^2)^{1/2}$  has a maximum at  $x = 0$  on  $[0, 1]$ . To find a function that looks like  $(1 - x^2)^{1/2}$  near  $x = 0$ , we could expand  $(1 - x^2)^{1/2}$  in a Taylor series around zero (i.e. the maximum).

In this example, let  $g(x) = (1 - x^2)^{1/2}$  and  $p(x) = 1$   
 $\implies g(x)p(x) = g(x)$  and

$$g(0) = 1, \quad g'(0) = 0, \quad g''(0) = -1$$

implying that

$$g(x) \approx 1 - \frac{1}{2}x^2, \quad \int_0^1 (1 - \frac{1}{2}x^2) dx = 1 - 1/6 = 5/6.$$

So, we could form

$$\tilde{p}(x) = \frac{6}{5} \cdot (1 - \frac{1}{2}x^2), \quad 0 < x < 1.$$

If we sample  $x_1, x_2, \dots, x_N$  i.i.d. from this  $\tilde{p}(\cdot)$  and use this MC estimator:

$$\hat{G}_{\tilde{p}, N} = \frac{1}{N} \sum_{i=1}^N \frac{5}{6} \cdot \frac{(1 - x_i^2)^{1/2}}{1 - \frac{1}{2} x_i^2}$$

it turns out that  $\text{var}(\hat{G}_{\tilde{p}, N}) = 0.011/N$ , about 1/5 of  $\text{var}(\hat{G}_N)$ .

In general, this is not a great enough reduction to make finding  $\tilde{p}(\cdot)$  worth the effort. But, suppose now that we generalize  $1 - \frac{1}{2}x^2$  to  $1 - \beta x^2$  so that

$$\int_0^1 (1 - \beta x^2) dx = 1 - \frac{1}{3} \beta \quad \Longrightarrow \quad \tilde{p}(x) = \frac{1 - \beta x^2}{1 - \beta/3}, \quad 0 < x < 1.$$

We now look for  $\beta$  that minimizes the variance of

$$\frac{(1 - x^2)^{1/2}(1 - \beta/3)}{1 - \beta x^2}$$

where  $x$  is drawn from  $\tilde{p}(x) = (1 - \beta x^2)/(1 - \beta/3)$ ,  $0 < x < 1$ . It turns out that  $\beta = 0.74$  minimizes this variance and the end result is

$$\text{var}(\hat{G}_{\tilde{p}, N}) = \frac{0.0029}{N}$$

which is a significant (by an order of magnitude) improvement compared with  $\text{var}(\hat{G}_N)$ .

**Example 2:** Estimate  $P[2 < X]$  for  $X \sim \text{Cauchy}$ :

$$G = \int_2^{\infty} [\pi(1 + x^2)]^{-1} dx.$$

Note that, for  $2 < x$ ,  $1/(1 + x^2)$  looks much like  $1/x^2$  and

$$\int_2^{\infty} \frac{1}{x^2} dx = \frac{1}{2}.$$

Then, we might try

$$\tilde{p}(x) = \frac{2}{x^2}, \quad 2 < x$$

and estimate  $G$  using

$$\hat{G}_{\tilde{p}, N} = \frac{1}{N} \sum_{i=1}^N \frac{x_i^2}{2\pi(1 + x_i^2)}$$

where  $x_1, x_2, \dots, x_N$  are i.i.d.  $\tilde{p}(\cdot)$ . Note that for  $X$  following  $\tilde{p}(x) = \frac{2}{x^2}$ ,  $2 < x$ , we have

$$Y = \frac{2}{X} \sim \text{uniform}(0, 1)$$

$\implies$  we could just sample  $y_1, y_2, \dots, y_N$  from  $\text{uniform}(0, 1)$  and take

$$x_i = \frac{2}{y_i}, \quad i = 1, 2, \dots, N.$$

In this example, variance reduction was not our major problem — we focused on finding a density that is easy to sample from.

**Example 3:** Find an MC estimator of

$$G = \int_0^1 x^{-1/2}(1-x)^{-1/2} dx.$$

Note that the integrand is the kernel of a beta density with parameters  $1/2$  and  $1/2$ , see the table of distributions handed out in class. The integrand has singularities at the endpoints and a "standard" MC estimator based on sampling from  $\text{uniform}(0, 1)$  will have infinite variance!

Define

$$g(x) = \frac{1}{x^{1/2}(1-x)^{1/2}}$$

and observe that it has singularities at  $x = 0$  and  $x = 1$ . Note that

$$g(x) \approx \begin{cases} 1/x^{1/2} & \text{for } x \text{ near } 0 \\ 1/(1-x)^{1/2} & \text{for } x \text{ near } 1 \end{cases}$$

and choose

$$\tilde{p}(x) \propto \frac{1}{x^{1/2}} + \frac{1}{(1-x)^{1/2}}$$

which is the sum of the approximating functions near the singularities. Now

$$\int_0^1 \left[ \frac{1}{x^{1/2}} + \frac{1}{(1-x)^{1/2}} \right] dx = 4 \implies \tilde{p}(x) = \frac{1}{4x^{1/2}} + \frac{1}{4(1-x)^{1/2}}$$

yielding

$$\hat{G}_{\tilde{p},N} = \frac{1}{N} \sum_{i=1}^N \frac{4}{x_i^{1/2} + (1 - x_i)^{1/2}}.$$

All terms in the above expression are bounded by 4 so  $\text{var}(\hat{G}_{\tilde{p},N})$  exists!

The above three examples illustrate importance sampling for

- variance reduction (Example 1),
- finding a density to sample from (Example 2), and
- producing bounded  $g(x)p(x)/\tilde{p}(x)$  even if  $g(x)p(x)$  is not (Example 3).

# Summary of Importance Sampling

We wish to estimate

$$G = \int g(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}.$$

Suppose that we have an “importance” density  $\tilde{p}(\mathbf{x})$  such that it approximates  $|g(\mathbf{x})| p(\mathbf{x})$  well, the support of  $\tilde{p}(\mathbf{x})$  covers that of  $p(\mathbf{x})$ , and it is easy to sample from  $\tilde{p}(\mathbf{x})$ . Then, the observation that

$$G = \int g(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} = \int \frac{g(\mathbf{x}) p(\mathbf{x})}{\tilde{p}(\mathbf{x})} \tilde{p}(\mathbf{x}) d\mathbf{x}$$

suggests the following algorithm:

**1)** Draw

$$\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N \stackrel{\text{i.i.d.}}{\sim} \tilde{p}(\mathbf{x})$$

where  $\tilde{p}(\mathbf{x}) \equiv$  the importance distribution.

**2)** Compute the importance weights

$$\phi_n = \frac{p(\mathbf{x}_n)}{\tilde{p}(\mathbf{x}_n)}, \quad n = 1, 2, \dots, N.$$

3) Approximate  $G$  by either of the following estimators:

$$\hat{G}_{\text{WE}} = \frac{\sum_{n=1}^N g(\mathbf{x}_n) \phi_n}{\sum_{n=1}^N \phi_n} \quad (\text{weighted average})$$

or

$$\hat{G} = \sum_{n=1}^N g(\mathbf{x}_n) \phi_n \quad (\text{simple average, discussed earlier}).$$

Note that  $\hat{G}$  is an unbiased estimator of  $G$  whereas  $\hat{G}_{\text{WE}}$  is asymptotically consistent (not necessarily unbiased).

Why is  $\hat{G}_{\text{WE}}$  asymptotically consistent? Because

$$\frac{1}{N} \sum_{n=1}^N g(\mathbf{x}_n) \phi_n \xrightarrow{p} \mathbb{E}_{\tilde{p}}[g(\mathbf{X}) \phi(\mathbf{X})] = G$$

and

$$\frac{1}{N} \sum_{n=1}^N \phi_n \xrightarrow{p} \mathbb{E}_{\tilde{p}}[\phi(\mathbf{X})] = 1.$$

Although  $\hat{G}_{\text{WE}}$  is biased, it often has smaller mean-square error than  $\hat{G}$ .

Since both numerator and denominator of  $\hat{G}_{\text{WE}}$  involve  $\phi_n$ , we



only need to know

$$\phi(\mathbf{x}) = \frac{p(\mathbf{x})}{\tilde{p}(\mathbf{x})}$$

up to a proportionality constant! Therefore,  $\hat{G}_{\text{WE}}$  is great for Bayesian computations where  $p(\mathbf{x})$  is often known only up to a proportionality constant.

# Sampling–Importance Resampling

Here, we use the importance sampling idea in a slightly different context.

Suppose we have  $N$  draws  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$  from a proposal distribution  $\tilde{p}(\mathbf{x})$ . Can we convert these samples to samples from a desired distribution  $p(\mathbf{x})$ ?

A *sampling–importance resampling method* for this conversion:

- For each  $\mathbf{x}_n$ , compute

$$\phi_n = \frac{p(\mathbf{x}_n)}{\tilde{p}(\mathbf{x})}$$
$$w_n = \frac{\phi_n}{\sum_{k=1}^N \phi_k}.$$

- Draw  $\mathbf{x}_\star$  from the discrete distribution over  $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$  with weight  $w_n$  on  $\theta_n$ .

If we need multiple samples  $\mathbf{x}_\star$ , it is suggested to draw them without replacement (which, of course, makes sense only if the number of samples to be drawn is a few times smaller than  $N$ ). If the number of samples to be drawn is  $N$ , then sample with replacement. The resampled  $\mathbf{x}_\star$  are drawn approximately from  $p(\mathbf{x})$ .

**Proof.** For simplicity, we focus on univariate  $x$ . Then

$$\begin{aligned}
 P\{x_* \leq a\} &= \overbrace{\sum_{n=1}^N w_n i_{-\infty, a}(x_n)}^{\text{empirical cdf}} \\
 &= \frac{N^{-1} \sum_{n=1}^N \phi_n i_{-\infty, a}(x_n)}{N^{-1} \sum_{n=1}^N \phi_n} \\
 &\rightarrow \frac{\mathbb{E}_{\tilde{p}}\left[\frac{p(X)}{\tilde{p}(X)} i_{-\infty, a}(X)\right]}{\mathbb{E}_{\tilde{p}}\left[\frac{p(X)}{\tilde{p}(X)}\right]} \\
 &= \frac{\int_{-\infty}^a p(x) dx}{\int_{-\infty}^{\infty} p(x) dx} = \int_{-\infty}^a p(x) dx.
 \end{aligned}$$

□

# MC Variance Reduction: Rao-Blackwellization

Suppose that we have drawn i.i.d. samples  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$  from  $p(\mathbf{X})$  and wish to estimate

$$I = \mathbb{E}_{p(\mathbf{x})}[g(\mathbf{X})].$$

Recall the straightforward estimator that we mentioned before:

$$\hat{G}_N = \frac{1}{N} \cdot [g(\mathbf{x}_1) + g(\mathbf{x}_2) + \dots + g(\mathbf{x}_N)]$$

known as the *simple average estimator* or *histogram estimator*.

Suppose that the random vector  $\mathbf{X}$  can be divided into two blocks:

$$\mathbf{X} = (\mathbf{X}^{(1)}, \mathbf{X}^{(2)})$$

where we can compute

$$\mathbb{E}_{p(\mathbf{x} | \mathbf{x}^{(2)})}[g(\mathbf{X}) | \mathbf{X}^{(2)} = \mathbf{x}^{(2)}]$$

analytically. Since the law of iterated expectations states

$$\mathbb{E}_{p(\mathbf{x}^{(2)})}[\mathbb{E}_{p(\mathbf{x} | \mathbf{x}^{(2)})}[g(\mathbf{X}) | \mathbf{X}^{(2)} = \mathbf{x}^{(2)}]] = \mathbb{E}_{p(\mathbf{x})}[g(\mathbf{X})]$$

or, more informally, after removing the annoying subscripts:

$$\mathbb{E} \{ \mathbb{E} [g(\mathbf{X}) | \mathbf{X}^{(2)} = \mathbf{x}^{(2)}] \} = \mathbb{E} [g(\mathbf{X})]$$

the following *Rao-Blackwellized estimator of G*

$$\begin{aligned}\tilde{G}_N &= \frac{1}{N} \cdot \left\{ \mathbb{E}_{p(\mathbf{x} | \mathbf{x}^{(2)})} [g(\mathbf{x}_1)] + \mathbb{E}_{p(\mathbf{x} | \mathbf{x}^{(2)})} [g(\mathbf{x}_2)] \right. \\ &\quad \left. + \cdots + \mathbb{E}_{p(\mathbf{x} | \mathbf{x}^{(2)})} [g(\mathbf{x}_N)] \right\}\end{aligned}$$

is unbiased. The above estimator is also known as the *mixture estimator*.

Which one of the above two MC estimators is better? Since

$$\begin{aligned}\text{var}[g(\mathbf{X})] &= \mathbb{E} \{ \text{var}[g(\mathbf{X}) | \mathbf{X}^{(2)} = \mathbf{x}^{(2)}] \} \\ &\quad + \text{var} \{ \mathbb{E} [g(\mathbf{X}) | \mathbf{X}^{(2)} = \mathbf{x}^{(2)}] \}\end{aligned}$$

we have

$$\text{var}[\hat{G}_N] = \frac{\text{var}[g(\mathbf{X})]}{N} \geq \frac{\text{var} \{ \mathbb{E} [g(\mathbf{X}) | \mathbf{X}^{(2)} = \mathbf{x}^{(2)}] \}}{N} = \text{var}[\tilde{G}_N].$$

Of course, the computational effort for obtaining the two estimates should also be taken into account when deciding which one is “better.”



A basic rule in MC computation: One should carry out analytical computations as much as possible!

## Comments:

- We can use the mixture estimator for density estimation as well: in particular, if we can compute  $p_{\mathbf{x}^{(1)} | \mathbf{x}^{(2)}}(\mathbf{x}^{(1)} | \mathbf{x}^{(2)})$  analytically, then

$$p_{\mathbf{x}^{(1)}}(\mathbf{x}^{(1)}) \approx \frac{1}{N} \cdot \left[ p_{\mathbf{x}^{(1)} | \mathbf{x}^{(2)}}(\mathbf{x}^{(1)} | \mathbf{x}_1^{(2)}) + p_{\mathbf{x}^{(1)} | \mathbf{x}^{(2)}}(\mathbf{x}^{(1)} | \mathbf{x}_2^{(2)}) + p_{\mathbf{x}^{(1)} | \mathbf{x}^{(2)}}(\mathbf{x}^{(1)} | \mathbf{x}_N^{(2)}) \right].$$

- If the samples  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$ , then it is not so straightforward to claim superiority of Rao-Blackwellization, but it has been shown for Gibbs sampler.

# ELEMENTARY TECHNIQUES FOR RANDOM SAMPLING

Our concern now is “sampling” observations from a given distribution using observations from another distribution [which is often  $\text{uniform}(0, 1)$ ].

## Inversion

**Probability Integral Transform:** If  $x$  is a continuous random variable with cdf  $F$ , then

$$F(x) \sim \text{uniform}(0, 1).$$

To use this idea for sampling from both continuous and discrete distributions, we extend the probability integral transform as follows.

**Theorem 1.** *Assume that we wish to sample a random variable  $x$  with cdf  $F(x)$ . Define*

$$x(u) = \min\{\tilde{x} : F(\tilde{x}) \geq u\}$$

*where  $u \sim \text{uniform}(0, 1)$ . Then*

$$x(u) \sim F.$$

**Proof.** Ripley, p. 59.  $\square$

For continuous  $x$  and known  $F^{-1}$ :

1. Sample  $u_i \sim$  i.i.d. uniform(0, 1),
2.  $x_i = F^{-1}(u_i)$ .

Then  $x_1, x_2, \dots, x_N \sim$  i.i.d.  $F$ .

**Example:** Sample Poisson(4):

$$p(x) = \frac{1}{x!} 4^x e^{-4}$$

So, if we generate  $u_1 = 0.324$  from uniform(0, 1), then  $x_1 = 3$ .

$x$	$p(x)$	cdf $F(x) = P[X \leq x]$
0	0.018316	0.018316
1	0.073262	0.091578
2	0.146525	0.238103
3	0.195367	0.433470
...	...	...



# Some Useful Inversion Formulas

pdf: $\mathbf{p}(x)$	cdf: $F(x)$	$X = F^{-1}(U)$	simplified form
<b>Exponential(<math>\lambda</math>)</b>			
$\lambda e^{-\lambda x}, x \geq 0$	$1 - e^{-\lambda x}$	$-\frac{1}{\lambda} \log(1 - U)$	$-\frac{1}{\lambda} \log(U)$
<b>Laplace(<math>\lambda</math>)</b>			
$\frac{\lambda}{2} e^{-\lambda x }, x \in \mathbb{R}^1$	$\frac{1}{2} e^{-\lambda x}, \text{ if } x \leq 0$	$\frac{1}{\lambda} \log(2U), \text{ if } U \leq \frac{1}{2}$	$\frac{1}{\lambda} \log(2U)$
	$1 - \frac{1}{2} e^{-\lambda x}, \text{ if } x > 0$	$-\frac{1}{\lambda} \log(2(1 - U)), \text{ if } U > \frac{1}{2}$	$-\frac{1}{\lambda} \log(2U)$
<b>Cauchy(<math>\theta</math>)</b>			
$\frac{\theta}{\pi(x^2 + \theta^2)}, x \in \mathbb{R}^1$	$\frac{1}{2} + \frac{1}{\pi} \arctan\left(\frac{x}{\theta}\right)$	$\theta \tan\left(\pi\left(U - \frac{1}{2}\right)\right)$	$\theta \tan(\pi U)$
<b>Rayleigh(<math>\sigma</math>)</b>			
$\frac{x}{\sigma} e^{-\frac{x^2}{2\sigma^2}}, x \geq 0$	$1 - e^{-\frac{x^2}{2\sigma^2}}$	$\sigma \sqrt{-\log(1 - U)}$	$\sigma \sqrt{-\log(U)}$
<b>Pareto(<math>a, b</math>)</b>			
$\frac{ab^a}{x^{a+1}}, x \geq b > 0$	$1 - \left(\frac{b}{x}\right)^a$	$\frac{b}{(1-U)^{1/a}}$	$\frac{b}{U^{1/a}}$

## Composition of Random Variables

The basic idea: be clever in manipulating functions that describe distributions. We now present one technique to illustrate the flavor of this approach. Consider sampling from a distribution that has the following form:

$$p(x) = \sum_{i=1}^K \alpha_i g_i(x), \quad \alpha_i > 0, \quad g_i(x) \geq 0.$$

Note that we do not require  $\int g_i(x) dx = 1$ .

In our original problem,  $g_i(x)$  may not be densities. However, we can write

$$p(x) = \sum_{i=1}^K \alpha_i \underbrace{\int g_i(x) dx}_{\beta_i} \cdot \frac{g_i(x)}{\underbrace{\int g_i(x) dx}_{h_i(x)}}.$$

Note that we need  $\sum \beta_i = 1$ , but this should hold if  $p(x)$  is a valid distribution.

We can find functions  $h_i(x)$  and coefficients  $\beta_i$  such that  $h_i(x) \geq 0$ ,  $\beta_i > 0$ ,  $\sum \beta_i = 1$ , and

$$\int h_i(x) dx = 1 \quad (\text{i.e. densities}).$$

## Sampling Scheme:

- We first sample a value  $m$  from the set  $\{1, 2, \dots, K\}$  with probabilities  $\beta_1, \beta_2, \dots, \beta_K$ ;
- Then, we take one observation from  $h_m(x)$ .

Repeating the above scheme  $N$  times, we end up with  $x_1, x_2, \dots, x_N$  from

$$h(x) = \sum_i \beta_i h_i(x) \quad (\text{a finite mixture}).$$

**Example:** We wish to sample from

$$p(x) = \frac{3}{5} + \frac{3}{5}x + \frac{3}{10}x^2, \quad 0 < x < 1.$$

Clearly, the above  $p(x)$  can be written as  $\sum_i \alpha_i g_i(x)$ . Over the interval  $(0, 1)$ , we know that  $h_1(x) = 1$ ,  $h_2(x) = 2x$  and  $h_3(x) = 3x^2$  are densities; hence

$$p(x) = \underbrace{\left(\frac{3}{5} \cdot 1\right)}_{\beta_1 h_1(x)} + \underbrace{\left(\frac{3}{10} \cdot 2x\right)}_{\beta_2 h_2(x)} + \underbrace{\left(\frac{1}{10} \cdot 3x^2\right)}_{\beta_3 h_3(x)} \quad \text{form } \sum_i \beta_i h_i(x).$$

Now, choose  $k = 1$  with probability  $\beta_1 = \frac{3}{5}$ ,  $k = 2$  with

probability  $\beta_2 = \frac{3}{10}$ , and  $k = 3$  with probability  $\beta_3 = \frac{1}{10}$  or

$$p(k) = \begin{cases} 6/10, & i = 1 \\ 3/10, & i = 2 \\ 1/10, & i = 3 \end{cases} .$$

We can use simple inversion to simulate samples from the above pmf.

### Entire algorithm:

1. Generate  $u_i \sim \text{uniform}(0, 1)$ .
2.
  - If  $u_i < 6/10$  set  $k = 1$ ;
  - If  $6/10 \leq u_i < 9/10$  set  $k = 2$ ;
  - If  $9/10 \leq u_i$  set  $k = 3$ .
3.
  - If  $k = 1$ , set  $x_i = w_1$  where  $w_1 \sim \text{uniform}(0, 1)$ ;
  - If  $k = 2$ , set  $x_i = \max\{w_1, w_2\}$   
where  $w_1, w_2 \sim \text{i.i.d. uniform}(0, 1)$ ;
  - If  $k = 3$ , set  $x_i = \max\{w_1, w_2, w_3\}$   
where  $w_1, w_2, w_3 \sim \text{i.i.d. uniform}(0, 1)$ .

### Note:

$$P[\max\{w_1, w_2\} < x] = x^2, \quad \text{for } x \geq 0$$
$$P[\max\{w_1, w_2, w_3\} < x] = x^3, \quad \text{for } x \geq 0.$$

## Grid Approach

Suppose that we wish to sample from  $p(x)$ . Here is how we can (approximately) accomplish that:

- Make a grid of values of  $x$  spanning the support of  $p(x)$ :

$$x_1, x_2, \dots, x_m.$$

For convenience, define also  $x_0 = -\infty$  and  $x_{m+1} = +\infty$ .

- Evaluate

$$p(x_1), p(x_2), \dots, p(x_m).$$

- Estimate the cdf of this distribution as follows:

$$\underbrace{0}_{\text{at } x_0}, \underbrace{0 + p(x_1)}_{\text{at } x_1}, \underbrace{0 + p(x_1) + p(x_2)}_{\text{at } x_2}, \dots, \underbrace{\sum_{i=1}^m p(x_i)}_{\text{at } x_m}, \underbrace{1}_{\text{at } x_{m+1}}.$$

- Generate a uniform(0, 1) random variable  $u$ .
- If  $u \in [\sum_{i=1}^{l-1} p(x_i), \sum_{i=1}^l p(x_i)]$ , draw  $x_l$ .

## Comments:

- **Nice properties:**

- It is easy to implement and understand (no tricks).

- **Disadvantages:**

- Draws are from an approximation to the true distribution.
- It is not clear how many grid points  $m$  should be chosen. If  $m$  is too small this will be a poor approximation. But how much is “too small”? Clearly, small is “good” in terms of computational complexity.
- Most values of a continuous random variable can not be generated by this scheme.

This approach has a larger educational (perhaps debugging?) than practical value.

## Basic Rejection Sampling

A powerful technique that allows us to sample from a distribution known only up to a constant. It is due to von Neumann, see

J. von Neumann, “Various techniques used in connection with random digits,” in *John von Neumann, Collected Works*, vol. V, A.H. Taub (Ed.), New York: Pergamon, 1961, pp. 768–770.

Also known as acceptance-rejection algorithm.

Say we wish to simulate from a distribution with density  $p(x)$ . In fact, rejection sampling “works” fine with discrete random variables and with random vectors (at least in principle — computational efficiency is important). Here, we focus on one-dimensional continuous random variables.

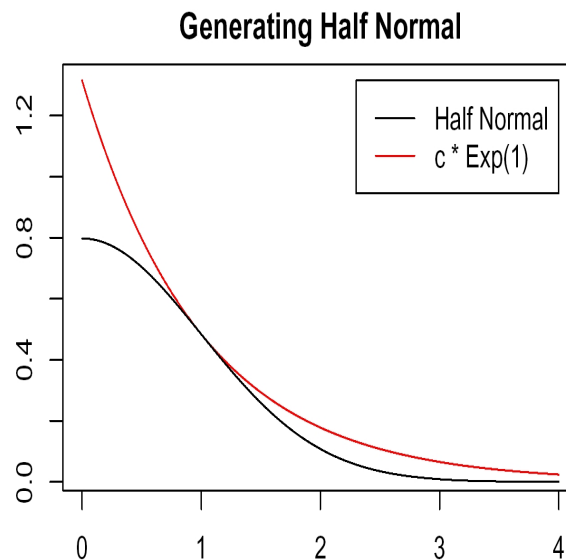
To implement this method, we need to find a dominating or “majorizing” pdf  $\tilde{p}(x)$  where

- $\tilde{p}(\cdot)$  is easy to sample from and

- 

$$p(x) \leq m \tilde{p}(x) = h(x)$$

for all  $x$  and some constant  $m > 1$ .



## Rejection Sampling Scheme:

1. Draw a proposal  $x$  from  $\tilde{p}(x)$  and compute the acceptance ratio:

$$r(x) = \frac{p(x)}{m \tilde{p}(x)} = \frac{p(x)}{h(x)} \leq 1.$$

2. Sample  $u \sim \text{uniform}(0, 1)$ .

- If  $u \leq r(x)$ , accept the draw and return  $x$ ;
- If  $u > r(x)$ , reject the draw and go back to 1 (and continue the loop).

Note that this step is equivalent to flipping a *biased coin* with success probability  $r(x)$ .

Then, the sample obtained using the above procedure is a draw from  $p(x)$ .



**Proof.** Let  $I$  be the indicator of whether a sample  $x$  is accepted. Then

$$\begin{aligned}
 P\{ \underbrace{I = 1}_{\text{sample } x \text{ accepted}} \} &= \int P\{I = 1|X = x\} \tilde{p}(x) dx \\
 &= \int r(x) \tilde{p}(x) dx \\
 &= \int \frac{p(x)}{m \tilde{p}(x)} \tilde{p}(x) dx = \frac{1}{m}.
 \end{aligned}$$

Next, we prove the desired result:

$$p(x|I = 1) = \frac{p(x)}{m \tilde{p}(x)} \cdot \tilde{p}(x) / P\{I = 1\} = \frac{p(x)}{m} \cdot m = p(x). \quad (2)$$

□

The acceptance probability is  $1/m$  and, clearly,  $m \geq 1$ . The number of trials until accepting a draw is a geometric random variable,  $\text{geometric}(1/m)$ ; hence, the average number of trials until acceptance is the mean of this geometric random variable, which is

$$\frac{1}{1/m} = m.$$

One consequence of this result is that  $m$  should be made as small as possible to minimize the number of rejections. The

optimal  $m$  is given by

$$m = \sup \frac{p(x)}{\tilde{p}(x)}.$$

Note that we do not need to find the best  $m$ , just one that satisfies

$$p(x) \leq m \tilde{p}(x) = h(x)$$

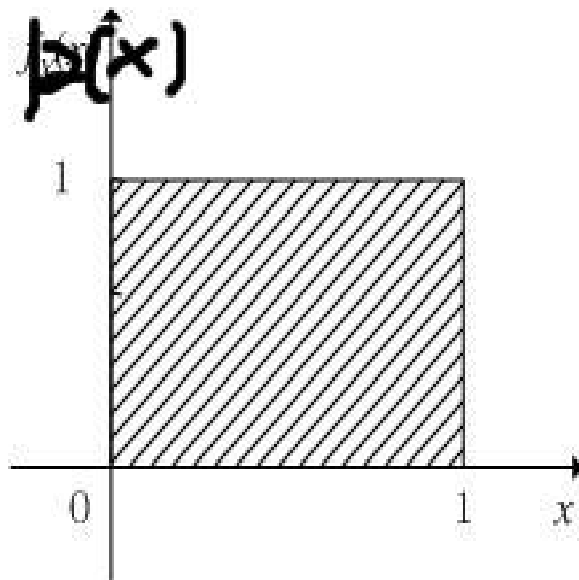
for all  $x$ .

## Comments:


- The key to rejection sampling is finding  $\tilde{p}(\cdot)$  with *correct support* such that  $\tilde{p}(\cdot)$  is *easy to sample from* and  $P[\text{accept } y]$  is *high*.
  - Intuitively,  $\tilde{p}(\cdot)$  needs to have *thicker tails* than  $p(\cdot)$  for  $p(x)/\tilde{p}(x)$  to remain bounded for all  $y$ . [For example, we cannot use a Gaussian pdf  $\tilde{p}(\cdot)$  to generate samples from a Cauchy pdf  $p(\cdot)$ . We can do the opposite, however.]
- Rejection sampling is self-monitoring — if the method is not working efficiently, very few draws will be accepted.

**Example 1 (trivial).**  $x \sim \text{uniform}(0, 1)$ :

$$p(x) = \begin{cases} 1, & 0 < x < 1, \\ 0 & \text{otherwise} \end{cases} .$$




We could sample by generating two uniform(0, 1) variables  $x$  and  $u$ , which we can visualize as



Accept  $x$  whenever  $u \leq 1$ . So, the candidate value is  $x$ ,  $m = 1$ ,  $p(x) = \text{uniform}(0, 1)$ ,  $\tilde{p}(x) = 1$ :

$$r(x) = \frac{p(x)}{m \tilde{p}(x)} = 1.$$

## Example:



$$p(x) = \frac{4}{\pi(1+x^2)}, \quad 0 < x < 1.$$

Here, we can still use  $\tilde{p}(x) = 1, 0 < x < 1$  because it has the correct support.

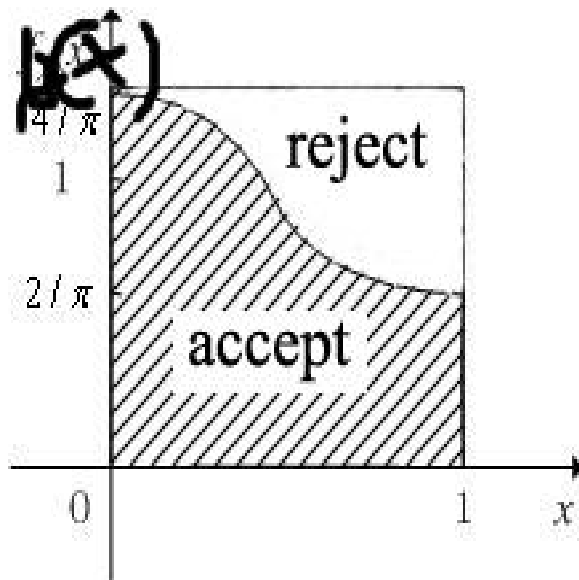
**Note:**  $m = 1$  (say) will not give the correct rejection region. Recall that we need

$$p(x) \leq m \tilde{p}(x)$$

which, in this case, will be satisfied if

$$m = \frac{4}{\pi}.$$

We sample uniformly over the region  $[0, 1] \times [0, 4/\pi]$  and reject values in the upper right corner:



For this example, we can write the general algorithm as follows:

1. Draw  $x$  from uniform  $(0, 1)$ .


$$\tilde{p}(y)$$

2. Draw  $u$  from uniform $(0, 1)$ .

3. If

$$u \leq r(x) = \frac{p(x)}{m \tilde{p}(x)} = \frac{1}{1 + x^2}$$

then return  $x$ . Otherwise go back to 1 and continue the loop.



### Example:

We wish to sample from a half-normal pdf with mean and variance parameters 0 and 1:

$$p(x) = \sqrt{2/\pi} \cdot \exp(-x^2/2) \cdot i_{[0,\infty)}(x).$$

One possibility: draw  $x \sim \mathcal{N}(0, 1)$  and reject all  $x < 0 \implies$  inefficient since 50% of the draws end up being rejected.

Consider using the exponential(1) proposal distribution:

$$\tilde{p}(x) = \exp(-x) i_{[0,\infty)}(x)$$

We need to find  $m$  such that

$$p(x) \leq m \tilde{p}(x) \geq \iff \sqrt{2/\pi} \cdot \exp(-x^2/2) \leq \exp(-x).$$

Differentiate the log of the above expression and find the smallest  $m$  such that the above inequality is satisfied. As a result, we obtain  $x = 1$  and the bound

$$m = \sqrt{\frac{2}{\pi}} \cdot \exp\left(\frac{1}{2}\right) \approx 1.315$$

which is the optimal  $m$  that we can choose in this case, leading to the acceptance rate of 76%.

To summarize, here is the rejection sampler for this case:

1. Draw  $x$  from the exponential(1) pdf and compute the acceptance ratio:

$$r(x) = \frac{p(x)}{m \tilde{p}(x)} = \exp[-0.5(x - 1)^2].$$

2. Sample  $u \sim \text{uniform}(0, 1)$ .

- If  $u \leq r(x)$ , accept the draw and return  $x$ ;
- If  $u > r(x)$ , reject the draw and go back to 1 (and continue the loop).

Note that this example is somewhat artificial since we do not need rejection to sample from a half normal distribution. Half normal distribution is the distribution of the absolute value of a standard normal random variable.

In the previous discussion, it was assumed that  $p(x)$  was a density function. In fact,  $p(x)$  only needs to be known up to a *multiplicative constant*:

$$l(x) = b p(x).$$

Here,  $b$  is the multiplicative constant that may be unknown. This case is particularly common in Bayesian inference where the posterior density is usually known up to a proportionality factor:

$$p(\theta|\text{data}) \propto p(\text{data}|\theta) \pi(\theta)$$

and the normalizing constant is difficult to calculate exactly.

Rejection sampling approach *does not* require knowing the constant  $b$ ! The original procedure can be modified as follows:

Find  $\tilde{m}$  such that

$$l(x) \leq \tilde{m} \tilde{p}(x) = h(x)$$

for all  $x$  and some positive constant  $\tilde{m}$ . Then, the rejection sampling scheme is:

1. Draw a proposal  $x$  from  $\tilde{p}(x)$  and compute the acceptance ratio:

$$r(x) = \frac{l(x)}{\tilde{m} \tilde{p}(x)} = \frac{l(x)}{h(x)} \leq 1.$$



2. Sample  $u \sim \text{uniform}(0, 1)$ .

- If  $u \leq r(x)$ , accept the proposal and return  $x$ ;
- If  $u > r(x)$ , reject the proposal and go back to 1 (and continue the loop).

Everything is the same as before, except the unnormalized density  $l(x)$  is used instead of the normalized density  $p(x)$ .

The acceptance probability for this scheme is  $b/\tilde{m}$ .

### **A general comment about the choice of $\tilde{p}(\cdot)$ :**

A good choice  $\tilde{p}(x)$  will normally be “close to”  $p(x)$  — we wish to minimize the separation between the two densities. Often, a parametric family of candidates  $\tilde{p}(\cdot)$  is chosen and the member from the parametric family with the smallest  $m$  or  $\tilde{m}$  is determined and used.

# A Trick for Finding $m$ or $\tilde{m}$ when Sampling from a Posterior Distribution

Say we want to sample  $\theta$ s from

$$\overbrace{p(\cdot)}^{p(\cdot)} p(\boldsymbol{\theta}|\mathbf{x}) \propto \overbrace{\pi(\boldsymbol{\theta}) p(\mathbf{x}|\boldsymbol{\theta})}^{l(\cdot)}.$$

by generating samples from the prior  $\overbrace{\pi(\boldsymbol{\theta})}^{\tilde{p}(\boldsymbol{\theta})}$  and rejecting some of them. Hence, our basic rejection algorithm is

1. Generate  $\boldsymbol{\vartheta}$  from  $\overbrace{\pi(\boldsymbol{\vartheta})}^{\text{proposal dist.}}$ .
2. Generate  $u$  from uniform(0, 1).
3. Repeat Steps 1 and 2 until

$$u \leq r(\boldsymbol{\vartheta}) = \frac{p(\mathbf{x}|\boldsymbol{\vartheta}) \pi(\boldsymbol{\vartheta})}{\tilde{m} \pi(\boldsymbol{\vartheta})} = \frac{p(\mathbf{x}|\boldsymbol{\vartheta})}{\tilde{m}}.$$

4. Return  $\boldsymbol{\theta} = \boldsymbol{\vartheta}$ .

Here, we need to find  $\tilde{m}$  such that

$$p(\mathbf{x}|\boldsymbol{\vartheta}) \pi(\boldsymbol{\vartheta}) \leq \tilde{m} \pi(\boldsymbol{\vartheta}) \iff p(\mathbf{x}|\boldsymbol{\vartheta}) \leq \tilde{m}$$

for all  $\vartheta$ . Choose  $\tilde{m} = \max_{\vartheta} p(\mathbf{x}|\vartheta)$ , which is the *maximized likelihood*!

This yields the following rejection-sampling scheme:

1. Generate  $\vartheta$  from  $\pi(\vartheta)$ .
2. Generate  $u$  from  $\text{uniform}(0, 1)$ .
3. Repeat Steps 1 and 2 until

$$u \leq \frac{p(\mathbf{x}|\vartheta)}{\max_{\vartheta} p(\mathbf{x}|\vartheta)}.$$

4. Return  $\theta = \vartheta$ .

Hence, those  $\vartheta$  from the prior  $\pi(\cdot)$  that are likely according to the likelihood are kept in the posterior sample! For example, a random draw equal to the ML estimate of  $\theta$

$$\hat{\theta} = \arg \max_{\theta} p(\mathbf{x}|\vartheta)$$

will *always* be accepted!

# Background: (Univariate) Slice Sampler

Consider now sampling a random variable  $\phi$  from a nonstandard  $p(\phi) \propto h(\phi)$ .

## (Seemingly Counter-Intuitive!) Idea:


- Invent a convenient bivariate distribution for, say,  $\phi$  and  $u$ , with marginal pdf for  $\phi$  specified by  $h(\phi)$ .
- Then, use Gibbs sampling to make

$$(\phi^{(0)}, u^{(0)}), (\phi^{(1)}, u^{(1)}), (\phi^{(2)}, u^{(2)}), \dots, (\phi^{(T)}, u^{(T)}).$$

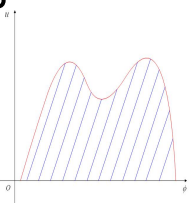


Create an auxiliary variable  $u$  just for convenience!


# (Univariate) Slice Sampler



“Invent” a joint distribution for  $\phi$  and  $u$  by declaring it to be

uniform on  :

$$p(\phi, u) = \begin{cases} \frac{1}{c}, & 0 < u < h(\phi) \\ 0, & \text{otherwise} \end{cases} \propto i_{(0, h(\phi))}(u).$$



With this joint pdf,  $P[\phi \leq 13] = \int_{-\infty}^{13} \frac{h(\phi)}{c} d\phi$ .




The marginal pdf of  $\phi$  is indeed specified by  $h(\phi) \implies$  if we figure out how to do Gibbs sampling, we know how to generate a  $\phi$  from  $h(\phi)$ .

# Gibbs Sampler is Easy in This Case!


$$p(u | \phi) = \text{uniform}(0, h(\phi))$$

$$p(\phi | u) = \text{uniform on } \underbrace{\{\phi | h(\phi) > u\}}_{\text{"slice"}}$$

**Step 1: Given  $\phi^{(t-1)}$ , sample  $u^{(t)} \sim \text{uniform}(0, h(\phi^{(t-1)}))$**




**Step 2: Given  $u^{(t)}$ , sample  $\phi^{(t)}$  Uniform from slice<sup>(t)</sup>**



If we can algebraically solve  $h(\phi) = u^{(t)}$ , our task is easy. What if not?

## Step 2 implementation using the rejection method

When we have band bounds on  $\phi$ , say  $\phi_{\text{MIN}} \leq \phi \leq \phi_{\text{MAX}}$



generate i.i.d. values  $\phi$  from  $\text{uniform}(\phi_{\text{MIN}}, \phi_{\text{MAX}})$  until we produce a  $\phi$  in the slice [i.e.  $h(\phi) > u^{(t)}$ ], which we then accept as  $\phi^{(t)}$ .

**Note:** For multivariate extensions of the slice sampler (particularly the “shrinkage idea”), see

R.M. Neal, “Slice sampling,” *Ann. Statist.*, vol. 31, pp. 705–741, June 2003.