# Tutorial on Monte Carlo Techniques 

Gabriel A. Terejanu<br>Department of Computer Science and Engineering<br>University at Buffalo, Buffalo, NY 14260<br>terejanu@buffalo.edu

## 1 Introduction

Monte Carlo (MC) technique is a numerical method that makes use of random numbers to solve mathematical problems for which an analytical solution is not known. The first article, "The Monte Carlo Method" by Metropolis and Ulam, has appeared for the first time in 1949 [9], even though well before that certain statistical problems were solved using random numbers. Since the simulation of random numbers is very time consuming, MC has became practical only with the advent of computers.

A simple MC simulation is the determination of $\pi$. Suppose we have a circle with radius $r=1$ inscribed within a square. Then the area ratio is:

$$
\begin{equation*}
\frac{A_{\text {circle }}}{A_{\text {square }}}=\frac{\pi r^{2}}{4 r^{2}}=\frac{\pi}{4} \tag{1}
\end{equation*}
$$

To determine $\pi$ we randomly select $n$ points, $p_{i}=\left(x_{i}, y_{i}\right)$ for $i=1 \ldots n$, in the square and approximate the area ratio by:

$$
\begin{equation*}
\frac{A_{\text {circle }}}{A_{\text {square }}}=\frac{m}{n} \tag{2}
\end{equation*}
$$

where $m$ is the number of random points in the circle, they must satisfy $x_{i}^{2}+y_{i}^{2} \leq 1$. Hence $\pi=4 \times \frac{m}{n}$.


Figure 1: Compute $\pi$
This method is called hit-or-miss Monte Carlo since the estimate is computed as the actual ratio of hits to random tries. It is the least efficient MC method.

## 2 Random Variables

The "Monte Carlo" name is derived from the city, with the same name, in the Principality of Monaco, well known for its casinos. This was because the roulette wheel was the simplest mechanical device for generating random numbers [10].

Even though "random variable" implies that one cannot predict its value, the distribution may well be known. The distribution of a random variable gives the probability of a given value [7].

The probability distribution of a discrete random variable is a list of probabilities associated with each of its possible values. It is also sometimes called the probability function or the probability mass function. (Valerie J. Easton and John H. McColl's Statistics Glossary v1.1).

A continuous random variable, $x$, takes any values in a certain interval $(a, b)$. It is defined by probability density function (pdf) $p(x)$ and the given interval. The probability of $x_{i}$ falling in an arbitrary interval $\left(a^{\prime}, b^{\prime}\right)$ is given by:

$$
\begin{equation*}
\mathrm{P}\left\{a^{\prime} \leq x \leq b^{\prime}\right\}=\int_{a^{\prime}}^{b^{\prime}} p(x) d x \tag{3}
\end{equation*}
$$

The pdf has to satisfy the following two conditions:

1. $p(x) \geq 0$ for any $x \in(a, b)$
2. $\int_{a}^{b} p(x) d x=1$

The expected value or mean, the second moment and the variance or central second moment of the random variable are respectively defined by:

$$
\begin{align*}
E[x]=\mu & =\int_{a}^{b} x p(x) d x  \tag{4}\\
E\left[x^{2}\right] & =\int_{a}^{b} x^{2} p(x) d x  \tag{5}\\
\operatorname{Var}[x]=\sigma^{2}=E\left[(x-\mu)^{2}\right] & =\int_{a}^{b}(x-\mu)^{2} p(x) d x=E\left[x^{2}\right]-\mu^{2} \tag{6}
\end{align*}
$$

where $\sigma$ is the standard deviation.
Consider now two continuous random variables $x$ and $y$. We say that $x$ and $y$ are statistically independent if the distribution of $x$ does not depend on $y$ and vice-versa. Hence the joint probability density function $f_{x y}(x, y)=f_{x}(x) f_{y}(y)$.

The covariance and the correlation of two random variables:

$$
\begin{align*}
\operatorname{Cov}[x, y] & =E[(x-E[x])(y-E[y])]=E[x y]-E[x] E[y]  \tag{7}\\
\operatorname{Corr}[x, y] & =\frac{\operatorname{Cov}[x, y]}{\sqrt{\operatorname{Var}[x] \operatorname{Var}[y]}} \tag{8}
\end{align*}
$$

Note: If $x$ and $y$ are uncorrelated then their covariance and the correlation coefficient is zero, hence $E[x y]=E[x] E[y]$. "Statistically independent random variables are always uncorrelated, but uncorrelated random variables can be dependent. Let $x$ be uniformly distributed over $[-1,1]$ and let $y=x^{2}$. The two random variables are uncorrelated but are clearly not independent" [8].

## 3 Generate Random Variables

The "Monte Carlo" name is derived from the city, with the same name, in the Principality of Monaco, well known for its casinos. This was because the roulette wheel was the simplest mechanical device for generating random numbers [10].
"A sequence of truly random numbers is unpredictable and therefore unreproducible. Such a sequence can only be generated by a random physical process" [7]. Practically it is very difficult to construct such a physical generator, which has to be fast enough, and to connect it to a computer. To overcome this problem one can use pseudo-random numbers, which are computed according to a mathematical formulation, hence are reproducible and appear random to someone who does not know the algorithm [12].

John von Neumann has constructed the first pseudorandom generator called mid-square method. Suppose we have a 4 -digit number $x_{1}=0.9876$. We square it, $x_{1}^{2}=0.97535376$, obtaining this way a 8 -digit number. We obtain $x_{2}$ by taking out the middle four digits, hence $x_{2}=0.5353$. Now square $x_{2}$ and so on. Unfortunately, the algorithm tends to produce disproportionate frequency of small numbers [10].

One popular algorithm is the multiplicative congruential method suggested by D.H. Lehmer in 1949. Given a modulus $m$, a multiplier $a$ and a starting point $x_{1}$, the method generates successive pseudo-random numbers by the formula [7]:

$$
\begin{equation*}
x_{i}=a x_{i-1} \operatorname{Mod}(m) \tag{9}
\end{equation*}
$$

## 4 MC integration

Let $f(x)$ be an arbitrary continuous function and $y=f(x)$ the new random variable. Then the expected value and the variance of $y$ :

$$
\begin{align*}
E[y]=E[f(x)] & =\int_{a}^{b} f(x) p(x) d x  \tag{10}\\
\operatorname{Var}[y]=\operatorname{Var}[f(x)] & =\int_{a}^{b}(f(x)-E[f(x)])^{2} p(x) d x \tag{11}
\end{align*}
$$

Note: $E[f(x)] \neq f(E[x])$.
Our goal is to calculate the expectation of $f(x)$ without computing the integral. This can be achieved by using a MC simulation.

Crude Monte Carlo: A simple estimate of the integral (10) can be obtained by generating $n$ samples $x_{i} \sim q(x)$, for $i=1 \ldots n$, and computing the estimate $I=\frac{1}{n} \sum_{i=1}^{n} f\left(x_{i}\right)$. To determine the accuracy of method we first have to introduce two theorems.

The law of large numbers implies that the average of a sequence of random variables, of a known distribution, converge to its expectation as the numbers of samples goes to infinity. Let us select $n$ numbers $x_{i}$ randomly with probability density $p(x)$. Hence:

$$
\begin{equation*}
I=\frac{1}{n} \sum_{i=1}^{n} f\left(x_{i}\right) \rightarrow E[f(x)]=\int_{a}^{b} f(x) p(x) d x \tag{12}
\end{equation*}
$$

In the statistical context $A(n)$ is a consistent estimator of $B$ if " $A(n)$ is said to converge to $B$ as $n$ goes to infinity if for any probability $P[0<P<1]$, and any positive quantity $\delta$, a $k$ can be found such that for all $n>k$ the probability that $A(n)$ will be within $\delta$ of $B$ is greater that $P$ " [7].

The central limit theorem: the sum of a large number of independent random variables is approximately normally distributed. A sum of identical and independent distribution (i.i.d) random variables $z_{i}$ with mean $\mu$ an finite variance $\sigma^{2}$ :

$$
\begin{equation*}
\frac{\sum_{i=1}^{n} z_{i}-n \mu}{\sigma \sqrt{n}} \rightarrow \mathcal{N}(0,1) \text { as } n \rightarrow \infty \tag{13}
\end{equation*}
$$

The error of the MC method [2]:

$$
\begin{align*}
\text { error } & \equiv|I-E[f(x)]| \\
& =\left|\frac{1}{n} \sum_{i=1}^{n} f\left(x_{i}\right)-E[f(x)]\right| \\
& =\frac{\sqrt{\operatorname{Var}[f(x)]}}{\sqrt{n}}\left|\frac{\sum_{i=1}^{n} f\left(x_{i}\right)-n E[f(x)]}{\sqrt{\operatorname{Var}[f(x)]} \sqrt{n}}\right| \\
& =\frac{\sqrt{\operatorname{Var}[f(x)]}}{\sqrt{n}}|\mathcal{N}(0,1)| \tag{14}
\end{align*}
$$

## Mathematical properties [7]

1. If the variance of $f(x)$ is finite, the MC estimate is consistent
2. The MC estimate is asymptotically unbiased
3. The MC estimate is asymptotically normally distributed
4. The standard deviation of the MC estimate is given by $\sigma=\frac{\sqrt{\operatorname{Var}[f(x)]}}{\sqrt{n}}$

Also note that the MC methods are immune to the curse of dimensionality. The accuracy of the method can be improved by increasing the number of samples, $n$, but the convergence is very slow. A better way to improve the accuracy is to decrease the variance, $\operatorname{Var}[f(x)]$. This methods are called variance-reducing techniques.

## 5 Variance Reducing Techniques

### 5.1 Stratified Sampling

This technique divides the full integration space into subspaces, for which a MC integration is performed. The final result is the sum of all partial results. We divide the integration domain $R$ into $k$ regions $R_{i}$ for $i=1 \ldots k$.

$$
\begin{equation*}
E[f(x)]=\int_{R} f(x) p(x) d x=\sum_{j=1}^{k} \int_{R_{j}} f(x) p(x) d x \tag{15}
\end{equation*}
$$

The MC estimate of the expectation becomes:

$$
\begin{equation*}
I=\sum_{j=1}^{k} \frac{\operatorname{Vol}\left(R_{j}\right)}{n_{j}} \sum_{i=1}^{n_{j}} f\left(x_{i}\right) \tag{16}
\end{equation*}
$$

where $n_{j}$ is the number of points used for MC integration on $R_{j}$, and $\operatorname{Vol}\left(R_{j}\right)$ is the subdomain volume.
The MC variance:

$$
\begin{align*}
\sigma^{2} & =\sum_{j=1}^{k} \frac{\operatorname{Vol}^{2}\left(R_{j}\right)}{n_{j}} \operatorname{Var}_{R_{j}}[f(x)]  \tag{17}\\
\text { where } \quad \operatorname{Var}_{R_{j}}[f(x)] & =\frac{1}{\operatorname{Vol}\left(R_{j}\right)} \int_{R_{j}}\left(f(x)-\frac{1}{\operatorname{Vol}\left(R_{j}\right)} \int_{R_{j}} f(x) p(x) d x\right)^{2} p(x) d x \tag{18}
\end{align*}
$$

By selecting carefully the number of points this can lead to a lower variance compared with the crude MC. If not it can perform worse than the crude MC.

### 5.2 Importance Sampling

The pdf under the integral, $p(x)$, may not be the best pdf for MC integration. In this case we want to use a different and simpler pdf $q(x)$ from which we can draw the samples. $q(x)$ is called the importance density or the proposal distribution. Therefore we can write:

$$
\begin{align*}
E[y]=E[f(x)] & =\int_{a}^{b} f(x) p(x) d x  \tag{19}\\
& =\int_{a}^{b} \frac{f(x) p(x)}{q(x)} q(x) d x  \tag{20}\\
& =E\left[\frac{f(x) p(x)}{q(x)}\right] \tag{21}
\end{align*}
$$

By generating $n$ samples $x_{i} \sim q(x)$, for $i=1 \ldots n$, the estimate becomes:

$$
\begin{equation*}
I=\frac{1}{n} \sum_{i=1}^{n} \frac{f\left(x_{i}\right) p\left(x_{i}\right)}{q\left(x_{i}\right)}=\frac{1}{n} \sum_{i=1}^{n} W\left(x_{i}\right) f\left(x_{i}\right) \tag{22}
\end{equation*}
$$

where $W\left(x_{i}\right)=\frac{p\left(x_{i}\right)}{q\left(x_{i}\right)}$ are the importance weights. Since the normalizing factor $p\left(x_{i}\right)$ is not know, we have to normalize the weights such that $\sum_{i=1}^{n} W\left(x_{i}\right)=1$ :

$$
\begin{equation*}
I=\frac{\frac{1}{n} \sum_{i=1}^{n} W\left(x_{i}\right) f\left(x_{i}\right)}{\frac{1}{n} \sum_{i=1}^{n} W\left(x_{i}\right)}=\frac{1}{n} \sum_{i=1}^{n} W_{n}\left(x_{i}\right) f\left(x_{i}\right) \tag{23}
\end{equation*}
$$

where $W_{n}\left(x_{i}\right)=\frac{W\left(x_{i}\right)}{\sum_{i=1}^{n} W\left(x_{i}\right)}$ are the normalized importance weights. This estimate is biased but the bias vanishes asymptotically as $n \rightarrow \infty$ [4].

The variance of the importance sampler estimate is given by [4]:

$$
\begin{equation*}
\operatorname{Var}_{q}[f(x)]=\frac{1}{n} \int\left[\frac{(f(x) p(x))^{2}}{q(x)}\right] d x-\frac{E_{p}^{2}[f(x)]}{n} \tag{24}
\end{equation*}
$$

To reduce the variance $q(x)$ should be chosen to match the shape of $p(x)$ or $|f(x)| p(x)$ [4].

### 5.3 Control Variates

Suppose that we have a function $g(x)$ for which $E[g(x)]$ is known. Let:

$$
\begin{equation*}
h(x)=f(x)-c(g(x)-E[g(x)]) \tag{25}
\end{equation*}
$$

where $h(x)$ can have a smaller variance than $f(x)$. Our estimate is given now by:

$$
\begin{equation*}
I=\frac{1}{n} \sum_{i=1}^{n}\left(f\left(x_{i}\right)-c g\left(x_{i}\right)\right)+c E[g(x)] \tag{26}
\end{equation*}
$$

The variance of $h(x)$ :

$$
\begin{equation*}
\operatorname{Var}[h(x)]=\operatorname{Var}[f(x)]+c^{2} \operatorname{Var}[g(x)]-2 c \operatorname{Cov}[f(x), g(x)] \tag{27}
\end{equation*}
$$

where $\operatorname{Cov}[f(x), g(x)]=E[(f(x)-E[f(x)])(g(x)-E[g(x)])]$.
We want to minimize the variance $\operatorname{Var}[h(x)]$ with respect to $c$ :

$$
\begin{equation*}
\frac{\mathrm{d} \operatorname{Var}[h(x)]}{\mathrm{d} c}=2 c \operatorname{Var}[g(x)]-2 \operatorname{Cov}[f(x), g(x)] \tag{28}
\end{equation*}
$$

By setting this to zero, the optimal value for $c$ is then given by:

$$
\begin{equation*}
c=\frac{\operatorname{Cov}[f(x), g(x)]}{\operatorname{Var}[g(x)]} \tag{29}
\end{equation*}
$$

Then the variance $\operatorname{Var}[h(x)]$ becomes:

$$
\begin{equation*}
\left.\operatorname{Var}[h(x)]=\operatorname{Var}[f(x)]-\frac{\operatorname{Cov}^{2}[f(x), g(x)]}{\operatorname{Var}[g(x)]}=\operatorname{Var}[f(x)]\left(1-\operatorname{Corr}^{2}[f(x), g(x)]\right)\right) \tag{30}
\end{equation*}
$$

where $\operatorname{Corr}[f(x), g(x)]=\frac{\operatorname{Cov}[f(x), g(x)]}{\sqrt{\operatorname{Var}[f(x)]} \sqrt{\operatorname{Var}[g(x)]}}$ is the correlation coefficient. Hence the quality of the estimation depends on the correlation between $f(x)$ and $g(x)$.

### 5.4 Rejection Sampling (or Acceptance Rejection Sampling)

Suppose that we know an upper bound for the underlying pdf $p(x)$ and we use a proposal distribution $q(x)$ as in importance sampling. So, there is $c<\infty$ such that $p(x)<c q(x)$ [4].

```
Algorithm 1 Rejection Sampling
    Select a uniform random variable \(u \sim \mathcal{U}(0,1)\)
    Draw a sample \(x \sim q(x)\)
    if \(u<\frac{p(x)}{c q(x)}\) then
        accept \(x\)
    else
        reject it and repeat the process
    end if
```

One should carefully select the value for $c$. If it is too small the rejection rate will be low; if it is too big the acceptance rate will be low.

### 5.5 Antithetic Variates

Consider that we have an even number of samples, $2 n$, drawn from $p(x)$. One approach is to generate correlated samples to reduce the variance by cancellations in their sum. The estimate:

$$
\begin{align*}
I & =\frac{1}{2 n} \sum_{i=1}^{2 n} f\left(x_{i}\right)  \tag{31}\\
& =\frac{1}{n} \sum_{i=1}^{n} \frac{f\left(x_{2 i-1}\right)+f\left(x_{2 i}\right)}{2}  \tag{32}\\
& =\frac{1}{n} \sum_{i=1}^{n} g\left(x_{2 i}\right) \tag{33}
\end{align*}
$$

where $g\left(x_{2 i}\right)=\frac{f\left(x_{2 i-1}\right)+f\left(x_{2 i}\right)}{2}$.
The variance $\operatorname{Var}\left[g\left(x_{2 i}\right)\right]$ :

$$
\begin{align*}
\operatorname{Var}\left[g\left(x_{2 i}\right)\right] & =\operatorname{Var}\left[\frac{f\left(x_{2 i-1}\right)+f\left(x_{2 i}\right)}{2}\right]  \tag{34}\\
& =\frac{1}{4}\left(\operatorname{Var}\left[f\left(x_{2 i-1}\right)\right]+\operatorname{Var}\left[f\left(x_{2 i}\right)\right]+2 \operatorname{Cov}\left[f\left(x_{2 i-1}\right), f\left(x_{2 i}\right)\right]\right) \tag{35}
\end{align*}
$$

Variance analysis:

1. If $\operatorname{Cov}\left[f\left(x_{2 i-1}\right), f\left(x_{2 i}\right)\right]=0$ or $f\left(x_{2 i-1}\right), f\left(x_{2 i}\right)$ are independent then $\operatorname{Var}\left[g\left(x_{2 i}\right)\right]=\frac{1}{2} \operatorname{Var}\left[f\left(x_{2 i-1}\right)\right]$, hence the variance of the estimator remains the same.
2. If $\operatorname{Cov}\left[f\left(x_{2 i-1}\right), f\left(x_{2 i}\right)\right]<0$ then the estimate is improved.
3. If $\operatorname{Cov}\left[f\left(x_{2 i-1}\right), f\left(x_{2 i}\right)\right]>0$ the actual performance is worse.

### 5.6 Partial averaging (or Rao-Blackwellization)

The method is based on the idea that "one should carry out analytical computation as much as possible" [3]. So, if we can analytically compute the integral over a subspace, this yields the same expected value but with a smaller variance.

From the Law of Iterated Expectations we can compute the conditional expectation of $f(x)$ for some sigma algebra $\mathcal{F}$ :

$$
\begin{equation*}
E[f(x)]=E[E[f(x) \mid \mathcal{F}]] \tag{36}
\end{equation*}
$$

From the Law of Total Variance:

$$
\begin{equation*}
\operatorname{Var}[f(x)]=\operatorname{Var}[E[f(x) \mid \mathcal{F}]]+E[\operatorname{Var}[f(x) \mid \mathcal{F}]] \tag{37}
\end{equation*}
$$

Therefore $\operatorname{Var}[f(x)] \geq \operatorname{Var}[E[f(x) \mid \mathcal{F}]]$. In practice we can factor $f(x)$ such that some of the parts to be computed analytical and the rest can be solved by using Monte Carlo simulation. The method can be very problem dependent due to this factorization and sometimes may be more computational expensive than just simulating the entire system [3].

## 6 Markov Chain Monte Carlo

MCMC deals with generating samples from some complex probability distribution $p(x)$ that is difficult to directly sample from. The approach is to use a sequence of samples $\left\{x_{k}\right\}$ from a Markov chain, such that the sequence converges on $p(x)$ as $k \rightarrow \infty$.

In the following section only the case of discrete-time, continuous state space Markov process is discussed.

### 6.1 Markov chain

A Markov process is a stochastic process that has the Markov property. A stochastic process has the Markov property if the conditional probability distribution of the future states given all the past states depends only on the current state:

$$
\begin{equation*}
p\left(x_{k+1} \mid x_{0}, x_{1} \ldots x_{k}\right)=p\left(x_{k+1} \mid x_{k}\right) \tag{38}
\end{equation*}
$$

A Markov chain is a sequence of random variables $\left\{x_{k}\right\}$ generated by a Markov processes. The value of $x_{k}$ is called the state of the process at time $k$. The one step transition probability density (transition kernel) $p(y \mid x)$ is the probability density that the process moves from $x$ to $y$. The transition kernel must satisfy:

$$
\begin{equation*}
\int p(y \mid x) d y=1 \tag{39}
\end{equation*}
$$

Definition: The Markov chain is said to be homogeneous/stationary if the transition kernels do not depend on the time.

Define $p(x)$ the total probability density of $x$, and $p\left(x_{0}\right)$ the initial probability density of $x_{0}$. The total probability density of $y$ is given by the Chapman-Kolmogorov equation:

$$
\begin{equation*}
p(y)=\int p(y \mid x) p(x) d x \tag{40}
\end{equation*}
$$

Theorem: "If a Markov chain is ergodic, then there exists a unique steady state distribution $\pi$ independent of the initial state" [4].

$$
\begin{equation*}
\pi(y)=\int p(y \mid x) \pi(x) d x \tag{41}
\end{equation*}
$$

Reversible condition (detailed balance property): Guarantees the invariance of $\pi$ under the transition kernel:

$$
\begin{equation*}
p(x \mid y) \pi(y)=p(y \mid x) \pi(x), \forall x, y \tag{42}
\end{equation*}
$$

"The unconditional probability of moving $y$ to $x$ is equal to the unconditional probability of moving $x$ to $y$, where $x$ and $y$ are both generated from $\pi(\cdot) "[4]$.
"In MCMC sampling framework the samples are generated by a homogeneous, reversible, ergodic Markov chain with invariant distribution $\pi "$ [4].

### 6.2 Metropolis-Hastings Algorithm

Published in 1953 by Nicholas Metropolis and generalized in 1970 by W.K.Hastings, the algorithm draws samples from the target density $\pi(x)$, for which the normalizing factor might be unknown.

A Markov chain is generated with a proposal density (candidate target) $q(y \mid x)$, where $\int q(y \mid x) d x=$ 1. The proposal density generates a value for $y$ when the process is at $x$ [5]. If $q(y \mid x)$ satisfies the reversible condition then our search for a transition kernel is over.

Suppose, without loss of generality, that:

$$
\begin{equation*}
q(y \mid x) \pi(x)>q(x \mid y) \pi(y) \tag{43}
\end{equation*}
$$

which means that the process moves from $x$ to $y$ much often than from $y$ to $x$ [5]. To reduce the number of moves from $x$ to $y$ one can introduce the probability of move, $\alpha(y \mid x)$, such that the transition kernel from $x$ to $y$ is given by:

$$
\begin{equation*}
p_{M H}(y \mid x) \equiv q(y \mid x) \alpha(y \mid x) \tag{44}
\end{equation*}
$$

From (43) set $\alpha(x \mid y)=1$, because we need as many moves as we can get from $y$ to $x$. To determine $\alpha(y \mid x)$, one must require $p_{M H}(y \mid x)$ to have the detailed balance property:

$$
\begin{align*}
p_{M H}(y \mid x) \pi(x) & =p_{M H}(x \mid y) \pi(y)  \tag{45}\\
q(y \mid x) \alpha(y \mid x) \pi(x) & =q(x \mid y) \alpha(x \mid y) \pi(y) \\
& =q(x \mid y) \pi(y) \tag{46}
\end{align*}
$$

Thus the probability of move is given by:

$$
\begin{equation*}
\alpha(y \mid x)=\min \left[\frac{q(x \mid y) \pi(y)}{q(y \mid x) \pi(x)}, 1\right] \tag{47}
\end{equation*}
$$

```
Algorithm 2 Metropolis-Hasting
    Draw \(x_{0}\) from \(p\left(x_{0}\right)\)
    for \(i=0\) to \(N\) do
        Generate \(y \sim q\left(y \mid x_{i}\right)\)
        Generate \(u \sim \mathcal{U}(0,1)\)
        if \(u \leq \alpha\left(y \mid x_{i}\right)\) then
            \(x_{i+1}=y\)
        else
            \(x_{i+1}=x_{i}\)
        end if
    end for
    return \(\left\{x_{N b}, x_{N b+1}, \ldots x_{N}\right\}\)
```

The number of steps, $N$, determines the algorithm performance. We need a good amount of steps, denoted $N b$, for the chain to approaches stationarity (burn-in period). Another important factor for the efficiency is to have many samples or a high acceptance rate for the values of $y$. The draws are regarded as samples from $\pi(x)$ only after the burn-in period $N b$ [5].

### 6.3 Gibbs Sampling

The method was introduced by German and German in 1984 and named after the physicist J.W. Gibbs. The algorithm is a special case of Metropolis-Hastings for generating a sequence of samples from a joint distribution $p(\mathbf{x})$ hard to sample, of two or more random variables.

Let $\mathbf{x}$ be a $n$-dimensional random vector: $\mathbf{x}=\left[x_{1}, x_{2}, \ldots x_{n}\right]^{T}$. The Gibbs sampler uses only univariate conditional distributions, where all of the random variables but one are assigned fixed values.

```
Algorithm 3 Gibbs Sampling
    Draw \(\mathbf{x}^{0}\) from \(p\left(\mathbf{x}^{0}\right)\), results : \(\mathbf{x}^{0}=\left[x_{1}^{0}, x_{2}^{0}, \ldots x_{n}^{0}\right]^{T}\)
    for \(i=1\) to \(N\) do
        Draw \(x_{1}^{i} \sim p\left(x_{1} \mid x_{2}^{i-1}, x_{3}^{i-1}, \ldots x_{n}^{i-1}\right)\)
        Draw \(x_{n}^{i} \sim p\left(x_{n} \mid x_{1}^{i-1}, x_{2}^{i-1}, \ldots x_{n-1}^{i-1}\right)\)
        \(\mathbf{x}^{i}=\left[x_{1}^{i}, x_{2}^{i}, \ldots x_{n}^{i}\right]^{T}\)
    end for
    return \(\left\{\mathbf{x}_{N b}, \mathbf{x}_{N b+1}, \ldots \mathbf{x}_{N}\right\}\)
```

where $N b$ denotes the burn-in period.
"The Gibbs sampler is a nice solution to estimation of hierarchical or structured probabilistic model. Gibbs sampling can be viewed as a Metropolis method in which the proposal distribution is defined in terms of the conditional distributions of the joint distribution and every proposal is always accepted" [4].

## 7 Sequential Monte Carlo

In this section we focus our attention on sequential state estimation using sequential Monte Carlo (SMC). SMC is known also as bootstrap filtering, particle filtering, the condensation algorithm, interacting particle approximations and survival of the fittest [1].

### 7.1 Sequential Importance Sampling (SIS)

Consider the following nonlinear system, described by the difference equation and the observation model:

$$
\begin{align*}
\mathbf{x}_{k} & =\mathbf{f}\left(\mathbf{x}_{k-1}, \mathbf{w}_{k-1}\right)  \tag{48}\\
\mathbf{z}_{k} & =\mathbf{h}\left(\mathbf{x}_{k}, \mathbf{v}_{k-1}\right) \tag{49}
\end{align*}
$$

Denote by $\mathbf{Z}_{k}=\left\{\mathbf{z}_{i} \mid 1 \leq i \leq k\right\}$ the set of all observations up to time $k$, conditionally independent given the process with distribution $p\left(\mathbf{z}_{k} \mid \mathbf{x}_{k}\right)$. Also, assume that the state sequence $\mathbf{x}_{k}$ is an unobserved (hidden) Markov process with initial distribution $p\left(\mathbf{x}_{0}\right)$ and transition distribution $p\left(\mathbf{x}_{k+1} \mid \mathbf{x}_{k}\right)$ [6].

Our aim is to estimate recursively the posterior distribution $p\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right)$ and expectations of the form[6]:

$$
\begin{equation*}
E\left[f\left(\mathbf{x}_{k}, \mathbf{w}_{k}\right)\right]=\int f\left(\mathbf{x}_{k}, \mathbf{w}_{k}\right) p\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right) d \mathbf{x}_{k} \tag{50}
\end{equation*}
$$

Suppose that we cannot sample from the posterior distribution and use an importance sampling approach to sample from a proposal distribution $q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right)$. Hence:

$$
\begin{align*}
E\left[f\left(\mathbf{x}_{k}, \mathbf{w}_{k}\right)\right] & =\int f\left(\mathbf{x}_{k}, \mathbf{w}_{k}\right) \frac{p\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right)}{q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right)} q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right) d \mathbf{x}_{k}  \tag{51}\\
& =\int f\left(\mathbf{x}_{k}, \mathbf{w}_{k}\right) \frac{p\left(\mathbf{Z}_{k} \mid \mathbf{x}_{k}\right) p\left(\mathbf{x}_{k}\right)}{p\left(\mathbf{Z}_{k}\right) q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right)} q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right) d \mathbf{x}_{k}  \tag{52}\\
& =\int f\left(\mathbf{x}_{k}, \mathbf{w}_{k}\right) \frac{W_{k}\left(\mathbf{x}_{k}\right)}{p\left(\mathbf{Z}_{k}\right)} q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right) d \mathbf{x}_{k} \tag{53}
\end{align*}
$$

where $W_{k}\left(\mathbf{x}_{k}\right)=\frac{p\left(\mathbf{Z}_{k} \mid \mathbf{x}_{k}\right) p\left(\mathbf{x}_{k}\right)}{q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right)} \propto \frac{p\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right)}{q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right)}$ are the unnormalized importance weights. One can get rid of the unknown normalizing density $p\left(\mathbf{Z}_{k}\right)$ by deriving the expectation as:

$$
\begin{align*}
E\left[f\left(\mathbf{x}_{k}, \mathbf{w}_{k}\right)\right] & =\frac{\int f\left(\mathbf{x}_{k}, \mathbf{w}_{k}\right) W_{k}\left(\mathbf{x}_{k}\right) q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right) d \mathbf{x}_{k}}{\left.\int p\left(\mathbf{Z}_{k} \mid \mathbf{x}_{k}\right) p\left(\mathbf{x}_{k}\right) \frac{q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right)}{q\left(\mathbf{x}_{k}\right)} \mathbf{Z}_{k}\right)} d \mathbf{x}_{k}  \tag{54}\\
& =\frac{E_{q}\left[W_{k}\left(\mathbf{x}_{k}\right) f\left(\mathbf{x}_{k}, \mathbf{w}_{k}\right)\right]}{E_{q}\left[W_{k}\left(\mathbf{x}_{k}\right)\right]} \tag{55}
\end{align*}
$$

By drawing $N$ samples $\left\{\mathbf{x}_{k}^{i}\right\} \sim q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right), i=1 \ldots N$, one can approximate the expectations:

$$
\begin{align*}
E\left[f\left(\mathbf{x}_{k}, \mathbf{w}_{k}\right)\right] & \approx \frac{\frac{1}{N} \sum_{i=1}^{N} W_{k}\left(\mathbf{x}_{k}^{i}\right) f\left(\mathbf{x}_{k}^{i}, \mathbf{w}_{k}\right)}{\frac{1}{N} \sum_{i=1}^{N} W_{k}\left(\mathbf{x}_{k}^{i}\right)}  \tag{56}\\
& =\sum_{i=1}^{N} \bar{W}_{k}\left(\mathbf{x}_{k}^{i}\right) f\left(\mathbf{x}_{k}^{i}, \mathbf{w}_{k}\right) \tag{57}
\end{align*}
$$

where $\bar{W}_{k}\left(\mathbf{x}_{k}^{i}\right)=\frac{W_{k}\left(\mathbf{x}_{k}^{i}\right)}{\sum_{i=1}^{N} W_{k}\left(\mathbf{x}_{k}^{i}\right)}$ are the normalized weights.
Since it is very hard to find a good proposal distribution in a high dimensional space, one can sequentially construct it. Suppose we can factorize the proposal distribution:

$$
\begin{equation*}
q\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right)=q\left(\mathbf{x}_{k} \mid \mathbf{x}_{k-1}, \mathbf{Z}_{k}\right) q\left(\mathbf{x}_{k-1}, \mathbf{Z}_{k-1}\right) \tag{58}
\end{equation*}
$$

The posterior distribution can be factorized as:

$$
\begin{align*}
p\left(\mathbf{x}_{k} \mid \mathbf{Z}_{k}\right) & =p\left(\mathbf{x}_{k-1} \mid \mathbf{Z}_{k-1}\right) \frac{p\left(\mathbf{z}_{k} \mid \mathbf{x}_{k}\right) p\left(\mathbf{x}_{k} \mid \mathbf{x}_{k-1}\right)}{p\left(\mathbf{z}_{k} \mid \mathbf{Z}_{k-1}\right)}  \tag{59}\\
& \propto p\left(\mathbf{x}_{k-1} \mid \mathbf{Z}_{k-1}\right) p\left(\mathbf{z}_{k} \mid \mathbf{x}_{k}\right) p\left(\mathbf{x}_{k} \mid \mathbf{x}_{k-1}\right) \tag{60}
\end{align*}
$$

Hence we can recursively update the weights:

$$
\begin{align*}
W_{k}\left(\mathbf{x}_{k}^{i}\right) & \propto \frac{p\left(\mathbf{x}_{k}^{i} \mid \mathbf{Z}_{k}\right)}{q\left(\mathbf{x}_{k}^{i} \mid \mathbf{Z}_{k}\right)}  \tag{61}\\
& =\frac{p\left(\mathbf{x}_{k-1}^{i} \mid \mathbf{Z}_{k-1}\right) p\left(\mathbf{z}_{k} \mid \mathbf{x}_{k}^{i}\right) p\left(\mathbf{x}_{k}^{i} \mid \mathbf{x}_{k-1}^{i}\right)}{q\left(\mathbf{x}_{k}^{i} \mid \mathbf{x}_{k-1}^{i}, \mathbf{Z}_{k}\right) q\left(\mathbf{x}_{k-1}^{i}, \mathbf{Z}_{k-1}\right)}  \tag{62}\\
& \propto W_{k-1}\left(\mathbf{x}_{k-1}^{i}\right) \frac{p\left(\mathbf{z}_{k} \mid \mathbf{x}_{k}^{i}\right) p\left(\mathbf{x}_{k}^{i} \mid \mathbf{x}_{k-1}^{i}\right)}{q\left(\mathbf{x}_{k}^{i} \mid \mathbf{x}_{k-1}^{i}, \mathbf{Z}_{k}\right)} \tag{63}
\end{align*}
$$

```
Algorithm 4 Sequential Importance Sampling
    1: Draw \(\mathbf{x}_{k}^{i} \sim q\left(\mathbf{x}_{k} \mid \mathbf{x}_{k-1}^{i}, \mathbf{Z}_{k}\right)\) for \(i=1 \ldots N\)
    2: Compute the importance weights \(W_{k}^{i}=W_{k-1}^{i} \frac{p\left(\mathbf{z}_{k} \mid \mathbf{x}_{k}^{i}\right) p\left(\mathbf{x}_{k}^{i} \mid \mathbf{x}_{k-1}^{i}\right)}{q\left(\mathbf{x}_{k}^{i} \mid \mathbf{x}_{k-1}^{i}, \mathbf{Z}_{k}\right)}\)
    3: Normalize the importance weights \(\bar{W}_{k}^{i}=\frac{W_{k}^{i}}{\sum_{i=1}^{N} W_{k}^{i}}\)
```

Degeneracy problem: After a few iterations all but one particle will have negligible weights. Hence the algorithms allots time to update a large number of weights with no effect in our sampling.

Proposition (Kong et al. 1994): The unconditional variance (that is, when the observations are regarded as random) of the importance ratios increases over time, where the importance ratio is given by:

$$
\begin{equation*}
W_{k}\left(\mathbf{x}_{k}^{i}\right) \propto \frac{p\left(\mathbf{x}_{k}^{i} \mid \mathbf{Z}_{k}\right)}{q\left(\mathbf{x}_{k}^{i} \mid \mathbf{Z}_{k}\right)} \tag{64}
\end{equation*}
$$

This problem can be overcome by adding a resampling strategy to SIS algorithm.

### 7.2 Sequential Importance Resampling (SIR)

The resampling step eliminates the samples with low importance weights and multiplies the ones with high importance weights. "Many of the ideas on resampling have stemmed from the work of Efron 1982, Rubin 1988, and Smith and Gelfand 1992" [11].

To detect the weights degeneracy (or sample impoverishment), one can use the effective sample size $N_{\text {eff }}$ given by [4]:

$$
\begin{align*}
N_{e f f} & =\frac{N}{1+\operatorname{Var}_{q}\left[\bar{W}_{k}\left(\mathbf{x}_{k}\right)\right]}  \tag{65}\\
& =\frac{N}{E_{q}\left[\bar{W}_{k}\left(\mathbf{x}_{k}^{2}\right)\right]} \tag{66}
\end{align*}
$$

An estimate of the effective sample size:

$$
\begin{equation*}
\hat{N}_{e f f}=\frac{N}{\sum_{i=1}^{N}\left(\bar{W}_{k}^{i}\right)^{2}} \tag{67}
\end{equation*}
$$

If $\hat{N}_{e f f}$ is less than an given threshold then run the resampling step.
One approach is to draw $N$ particles from the current set of particles with probabilities proportional to their weights and set the new weights to $\frac{1}{N}$. This can be done by generating a new set of indices from a CDF that maps the random measure $\left\{\mathbf{x}_{k}^{i}, \bar{W}_{k}^{i}\right\}$ into an equally weighted random measure $\left\{\mathbf{x}_{k}^{j}, \frac{1}{N}\right\}$ [11].


Figure 2: Resampling process [11]

```
Algorithm 5 Resampling Algorithm
    Initialize CDF : \(c_{1}=0\)
    Construct CDF : \(c_{i}=c_{i-1}+\bar{W}_{k}^{i}\), for all \(i=1 \ldots N\)
    for \(j=1\) to \(N\) do
        Sampling index: \(u \sim \mathcal{U}(0,1)\)
        Resampling index: \(i=C D F^{-1}(u) N\)
        Assign sample: \(\tilde{\mathbf{x}}_{k}^{j}=\mathbf{x}_{k}^{i}\)
        Assign weight: \(\bar{W}_{k}^{j}=\frac{1}{N}\)
    end for
    return the new set \(\left\{\tilde{\mathbf{x}}_{k}^{j}, \bar{W}_{k}^{j}\right\}\)
```

where $C D F^{-1}(\cdot)$ is the inverse of the $C D F(\cdot)$ function.
This step reduces the sample impoverishment effect but introduces new practical problems: it limits the opportunity to have a parallel algorithm, particles with high weights are statistically selected many times and leads to a loss of diversity among the particles [1].

## References

[1] M.S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking. IEEE Trans. on Signal Processing, 50, 2002.
[2] Paul J. Atzberger. Strategies for Improving the Efficiency of Monte-Carlo Methods. http://www.cims.nyu.edu/ paulatz/class_finance/improvingMonteCarloMethods.pdf.
[3] K.R. Beevers. Monte Carlo integration. Lecture Notes for CSCI-6971, 2006.
[4] Zhe Chen. Bayesian Filtering: From Kalman Filters to Particle Filters, and Beyond. Manuscript.
[5] S. Chib and E. Greenberg. Understanding the Metropolis-Hastings algorithm. The American Statistician, 1995.
[6] A. Doucet, S. Godsill, and C. Andrieu. On Sequential Monte Carlo Sampling Methods for Bayesian Filtering. Technical report, Signal Processing Group, Department of Engineering, University of Cambridge, UK.
[7] F. James. Monte Carlo Theory and Practice. Rep. Prog. Phys., 43:1147-1188, 1980.
[8] D. Johnson. Jointly Distributed Random Variables. http://cnx.org/content/m11248/latest/, 2005.
[9] N. C. Metropolis and S. M. Ulam. The Monte-Carlo Method. J. Amer. Stat. Assoc., 44:335-341, 1949.
[10] Ilya M. Sobol'. A Primer for the Monte Carlo Method. CRC Press, 1994.
[11] Rudolph van der Merwe, Nando de Freitas, Arnaud Doucet, and Eric Wan. The unscented particel filter. In Advances in Neural Information Processing Systems 13, Nov 2001.
[12] S. Weinzierl. Introduction to Monte Carlo Methods. Technical report, NIKHEF Theory Group, 2000.

