# MARKOV CHAIN MONTE CARLO: THE METROPOLIS HASTINGS ALGORITHM

ELIOT HUANG

Abstract. This paper develops the mathematical foundations of Markov chains and applies them to the study of Markov Chain Monte Carlo (MCMC) methods, with a focus on the Metropolis–Hastings algorithm. Beginning with the definition of discrete-time Markov chains, we establish core properties including irreducibility, aperiodicity, recurrence, and the existence of stationary distributions. These results culminate in a proof of the ergodic theorem via a coupling argument, demonstrating the convergence of long-run empirical averages to expectations under the invariant distribution. Building upon these principles, we introduce and derive the Metropolis–Hastings algorithm as a canonical example of MCMC, highlighting its role in constructing reversible chains with prescribed stationary distributions.

Contents



## 1. Introduction

Markov chains form one of the most central objects in probability theory, capturing the behavior of systems that evolve stochastically in time with memory limited to their present state. Their elegant structure allows for deep theoretical analysis, while their ubiquity makes them indispensable tools across mathematics, statistics, and physics. One particularly important application is the simulation of complex distributions via Markov Chain Monte Carlo (MCMC), a family of algorithms that leverages the long-run behavior of chains to approximate otherwise intractable quantities.

Among MCMC methods, the Metropolis–Hastings algorithm stands as both foundational and versatile. By carefully designing transition rules that satisfy detailed balance with respect to a target distribution, the algorithm ensures reversibility and invariance, guaranteeing convergence under broad conditions. In practice,

*Date*: DEADLINES: Draft AUGUST 20 and Final version September 3, 2025.

this provides a systematic way to sample from complicated probability measures, enabling applications ranging from Bayesian statistics to statistical physics.

The aim of this paper is twofold. First, we provide a rigorous introduction to the theory of discrete-time Markov chains, including definitions, key properties, and convergence results. In particular, we establish the ergodic theorem for Markov chains via coupling, a constructive method that illuminates both the dynamics of convergence and the probabilistic structure underlying ergodicity. Second, we apply these principles to the Metropolis–Hastings algorithm, showing how its design reflects and exploits the fundamental properties of Markov chains. By bridging theoretical results with algorithmic construction, we hope to clarify both why the algorithm works and how it connects to broader themes in stochastic processes and probability theory.

## 2. Discrete-Time Markov Chains

**Definition 2.1** (Stochastic Process)**.** A **stochastic process** is a collection of random variables $X_t$ taking values in a state space $S$ and indexed by time. Within the context of this paper, time remains constrained to a subset $N \subseteq \{0, 1, 2, \dots\} = \mathbb{N}_0$. Hence, every stochastic process considered is a **discrete-time** process.

**Definition 2.2** (Markov Property)**.** A stochastic process $X_t$ with values in $S$ satisfies the **Markov property** if for all $n \geq 1$ and all states $i_0, \dots, i_n \in S$,

$$\mathbb{P}\{X_n = i_n \mid X_{n-1} = i_{n-1}, \dots, X_0 = i_0\} \ = \ \mathbb{P}\{X_n = i_n \mid X_{n-1} = i_{n-1}\}.$$

In other words, the future is independent of the past given the present.

**Definition 2.3** (Markov Chain)**.** A (discrete-time) **Markov chain** on a state space $S$ is a discrete-time stochastic process obeying the Markov property.

A Markov chain is considered discrete if $S$ is finite or countably infinite. In this paper, one may assume all Markov chains discussed are discrete.

**Definition 2.4** (Transition Probabilities and Transition Matrix)**.** The **transition probabilities** of a Markov chain are

$$p_t(i, j) = \mathbb{P}\{X_{t+1} = j \mid X_t = i\}.$$

In the finite case, the associated **transition matrix** (or stochastic matrix) is the matrix $\mathbf{P} = [p(i, j)]_{i,j \in S}$ with nonnegative entries and where each row sums to 1.

**Definition 2.5** (Time-Homogeneity)**.** A Markov chain is **time-homogeneous** if for all $t$ and $i, j \in S$,

$$p_t(i, j) = \mathbb{P}\{X_{t+1} = j \mid X_t = i\} = \mathbb{P}\{X_t = j \mid X_{t-1} = i\}.$$

We observe that $p_t(i, j)$ depends only on the states $i, j$ and not on $t$. Thus, for a time-homogeneous Markov chain, we may drop the subscript and simply write $p(i, j)$ for the transition probabilities.

For the remainder of the paper, unless stated otherwise, Markov chains can be assumed to be time-homogeneous.

**Definition 2.6** ($n$-step Transition Probabilities)**.** For $n \geq 1$ and $i, j \in S$, define

$$p^{(n)}(i, j) \ = \ \mathbb{P}\{X_n = j \mid X_0 = i\}.$$

Given the relevant Markov chain is time homogeneous, we observe by Definition 2.4,

$$\mathbb{P}\{X_{n+k} = j \mid X_k = i\} = p^{(n)}(i,j) \text{ for all } k \geq 0.$$

**Proposition 2.7** (Chapman–Kolmogorov Equations)**.** *For all $m, n \in \mathbb{N}$ and $i, j \in S$,*

$$p^{(m+n)}(i,j) \ = \ \sum_{k \in S} p^{(m)}(i,k)\, p^{(n)}(k,j).$$

*Equivalently,* $\mathbf{P}^{m+n} = \mathbf{P}^m \mathbf{P}^n$.

*Proof.* Suppose $i, j \in S$. Let $m, n \in \mathbb{N}$. By Definition 2.6,

$$\begin{aligned}
p^{(m+n)}(i,j) &= \mathbb{P}\{X_{m+n} = j \mid X_0 = i\} \\
&= \sum_{k \in S} \mathbb{P}\{X_{m+n} = j, X_n = k \mid X_0 = i\} \text{ (law of total probability)} \\
&= \sum_{k \in S} \mathbb{P}\{X_{m+n} = j \mid X_n = k, X_0 = i\} \mathbb{P}\{X_n = k \mid X_0 = i\} \\
&= \sum_{k \in S} \mathbb{P}\{X_{m+n} = j \mid X_n = k\} \mathbb{P}\{X_n = k \mid X_0 = i\} \text{ (Markov property)} \\
&= \sum_{k \in S} \mathbb{P}\{X_m = j \mid X_0 = k\} \mathbb{P}\{X_n = k \mid X_0 = i\} \text{ (Definition 2.4)} \\
&= \sum_{k \in S} p^{(m)}(k,j) p^{(n)}(i,k) \text{ (Definition 2.6).}
\end{aligned}$$

Thus, we conclude

$$p^{(m+n)}(i,j) = \sum_{k \in S} p^{(m)}(k,j)\, p^{(n)}(i,k),$$

as desired. □

**Remark 2.8.** As an immediate consequence of Proposition 2.7, for any $i, j, k \in S$ and $m, n \geq 0$,

$$p^{(m+n)}(i,j) = \sum_{k' \in S} p^{(m)}(i,k')\, p^{(n)}(k',j) \ \ \geq \ \ p^{(m)}(i,k)\, p^{(n)}(k,j),$$

since each term in the sum is nonnegative. In particular, if $p^{(m)}(i,k) > 0$ and $p^{(n)}(k,j) > 0$, then $p^{(m+n)}(i,j) > 0$.

**Definition 2.9** (Stationary Distribution)**.** Let $X_t$ be a Markov chain with transition kernel $\mathbf{P}$ on a state space $S$. A probability vector $\pi$ on $S$ is called a **stationary distribution** if

$$\pi = \pi \mathbf{P},$$

or equivalently

$$\pi(j) = \sum_{i \in S} \pi(i) \mathbf{P}(i,j) \quad \text{for all } j \in S.$$

Note, in the finite case, $\pi$ is a left eigenvector with eigenvalue 1.

Stationary distributions are also commonly known as invariant, equilibrium, or steady-state distributions.

**Definition 2.10** (Communication)**.** A state $j$ is **accessible** from $i$ if there exists $n \geq 0$ with $p^{(n)}(i,j) > 0$. We denote $j$ being accesible from $i$ as $i \to j$. States $i$ and $j$ **communicate** (denoted $i \leftrightarrow j$) if each is accessible from the other, i.e. there exist $m, n$ such that $p^{(m)}(i,j) > 0$ and $p^{(n)}(j,i) > 0$.

**Proposition 2.11.** *The communication relation $\leftrightarrow$ on the state space $S$ is an equivalence relation.*

*Proof.* Let $X_n$ be a Markov chain on a state space $S$. Consider states $i, j \in S$.

We observe if a Markov chain is in state $i$ then in 0 steps, the Markov chain must certainly remain in state $i$, and so $p_0(i,i) = 1$. Since $p_0(i,i) = 1 > 0$, $i \leftrightarrow i$. Thus, $\leftrightarrow$ is clearly reflexive.

Next, if $i$ communicates with $j$, then immediate from Definition 2.10, $j$ must communicate with $i$, and so $\leftrightarrow$ is symmetric.

Lastly, consider states $i, j, k \in S$, where $i \leftrightarrow j$ and $j \leftrightarrow k$. By Definition 2.9, there exist $m, n \geq 0$ such that $p^{(m)}(i,j) > 0$ and $p^{(n)}(j,k) > 0$. Thus, we can then demonstrate

$$\begin{aligned}
p^{(m+n)}(i,k) &= \mathbb{P}\{X_{m+n} \mid X_0 = i\} \\
&\geq \mathbb{P}\{X_{m+n}, X_m = j \mid X_0 = i\} \\
&= \mathbb{P}\{X_m = j \mid X_0 = i\}\mathbb{P}\{X_{m+n} = k \mid X_m = j\} \\
&= p^{(m)}(i,j)p^{(n)}(j,k) \\
&> 0.
\end{aligned}$$

By similar process, we can show that given there exist $m', n' \geq 0$ such that $p^{(m')}(j,i) > 0$ and $p^{(n')}(k,j) > 0$, $p^{(m'+n')}(k,i) > 0$. By Definition 2.9, $i \leftrightarrow k$, and therefore $\leftrightarrow$ is transitive. □

From the equivalence relation $\leftrightarrow$, the state space $S$ is uniquely partitioned into disjoint equivalence classes. The equivalence classes of the relation $\leftrightarrow$ are referred to as **communication classes**.

**Definition 2.12** (Irreducibility)**.** A Markov chain (or transition matrix $\mathbf{P}$) is known as **irreducible** if the state space $S$ is a single communication class.

In other words, a Markov chain is irreducible if $i \to j$ for every $i, j \in S$.

**Definition 2.13** (Period)**.** For a state $i$, define the **period** of $i$, $d = d(i)$, to be the greatest common divisor of the set

$$J_i = \{n \geq 0 \mid p^{(n)}(i,i) > 0\}.$$

**Proposition 2.14** (Period is a Class Property)**.** *If $i$ and $j$ communicate, then $d(i) = d(j)$. Consequently, in an irreducible chain, all states have the same period; therefore, we can say the chain (or $\mathbf{P}$) has a period $d$.*

*Proof.* Fix $i, j \in S$ such that $i$ and $j$ communicate. By Definition 2.10, there exist $m, n \geq 0$ such that

$$p^{(m)}(i,j) > 0 \quad \text{and} \quad p^{(n)}(j,i) > 0.$$

Let $r \in J_j$. By Definition 2.13, $p^{(r)}(j,j) > 0$. Then, by the Chapman–Kolmogorov equations (Proposition 2.7),

$$p^{(m+r+n)}(i,i) \geq p^{(m)}(i,j)p^{(r)}(j,j)p^{(n)}(i,j) > 0.$$

Thus, for all $r \in J_j$, $m+r+n \in J_i$. Therefore, $d(i)$ must divide $m+r+n$, so $k_1 d(i) = m+r+n$ for some integer $k_1$. Now, we observe since $p^{(m)}(i,j) > 0$ and $p^{(n)}(j,i) > 0$, from Remark 2.8, we can see that $p^{(m+n)}(i,i) > 0$ and $p^{(n+m)}(j,j) > 0$, and therefore $m+n \in J_J$ and $m+n \in J_i$. Critically, this means for some integer $k_2$, $k_2 d(i) = m+n$. We can easily show

$$r = (m+r+n) - (m+n) = k_1 d(i) - k_2 d(i) = (k_1 - k_2) d(i).$$

Since $k_1 - k_2$ is an integer, $r$ is divisible by $d(i)$. Thus, for all $r \in J_j$, $d(i)$ divides $r$, and so $d(i)$ divides $d(j)$. By symmetry, $d(j)$ divides $d(i)$, and so $d(i) = d(j)$. □

**Definition 2.15** (Aperiodicity)**.** We refer to an irreducible chain as **aperiodic** if $d = 1$.

**Definition 2.16** (Recurrence and Transience)**.** If a Markov chain is irreducible and finite, the chain is **recurrent**.

Now, consider the irreducible and countably infinite case. Let $T_i = \inf\{n \geq 1 \mid X_n = i\}$ represent the the first return time to the state $i$. An irreducible and countably infinite chain is **recurrent** if for every state $i \in S$,

$$\mathbb{P}\{T_i < \infty \mid X_0 = i\} = 1.$$

Simply put, a recurrent Markov chain returns infinitely often to each state.

If

$$\mathbb{P}\{T_i < \infty \mid X_0 = i\} < 1,$$

i.e. the chain is not recurrent, then the chain is **transient**. If the chain is transient, then each state almost surely receives only a finite amount of visits.

**Definition 2.17** (Positive and Null Recurrence)**.** A recurrent chain is **null recurrent** if

$$\lim_{n \to \infty} p^{(n)}(x,y) = 0.$$

A recurrent chain that is not null recurrent is called **positive recurrent**. Note, every finite irreducible Markov chain is positive recurrent.

**Theorem 2.18.** *Let* $\mathbf{P}$ *be irreducible. The following are equivalent:*

*(i) every state is positive recurrent,*
*(ii) some state* $i$ *is positive recurrent,*
*(iii)* $\mathbf{P}$ *has an invariant distribution* $\pi$*.*

*Proof.* We do not include a proof of this result here, as the proof of the result would take us too far afield. For the interested reader, a complete proof can be found in Markov chain notes offered by Cambridge University [6]. □

**Definition 2.19** (Ergodicity)**.** A Markov chain with countable $S$ that is irreducible, aperiodic, and positive recurrent is called **ergodic**.

As a consequence of Definition 2.17, a finite Markov chain is ergodic if the chain is irreducible and aperiodic.

**Lemma 2.20.** *Suppose* $\mathbf{P}$ *is the transition matrix of an irreducible, aperiodic Markov chain. Then for any* $i, j \in S$*, there exists* $N(i,j) > 0$ *such that*

$$p^{(n)}(i,j) > 0 \quad \textit{for all } n \geq N(i,j).$$

*Proof.* Suppose $\mathbf{P}$ is irreducible and aperiodic. Fix $i, j \in S$. Since $\mathbf{P}$ is irreducible, by Definition 2.12, there exists some $m(i,j) \geq 0$ such that $p^{(m(i,j))}(i,j) > 0$. Also, since $\mathbf{P}$ is aperiodic, for each state $j$ there exists an infinite set of return times $J_j = \{n \geq 1 \mid p^{(n)}(j,j) > 0\}$ with $d(j) = 1$. This implies that there exists $M(j) > 0$ such that for all $n \geq M(j)$, $n \in J_j$ and therefore $p^{(n)}(j,j) > 0$.

Now, for any $n \geq M(j)$ we can write

$$p^{(n+m(i,j))}(i,j) \;\geq\; p^{(m(i,j))}(i,j)\, p^{(n)}(j,j) \;>\; 0,$$

by Remark 2.8.

Therefore, if we define $N(i,j) = M(j) + m(i,j)$, then $p^{(n)}(i,j) > 0$ for all $n \geq N(i,j)$. ☐

**Theorem 2.21** (Ergodic Theorem for Markov Chains)**.** *If $X_n$ is a finite ergodic Markov chain with transition matrix $\mathbf{P}$, then there exists a unique stationary distribution $\pi$ such that*

$$\lim_{n\to\infty} \mathbf{P}^n = 1\pi,$$

*for any initial distribution.*

*Proof.* Immediately, by Definition 2.17 and Definition 2.19, we know an ergodic finite chain must be positive recurrent. Thus, by Theorem 2.18, there then exists an invariant distribution $\pi$ for $\mathbf{P}$.

Now, we want to show that $\mathbf{P}^n$ converges to $\pi$ regardless of initial distribution.

Let $X_n$ and $Y_n$ be independent Markov chains with initial distributions $\lambda$ and $\pi$ and a shared transition kernel $\mathbf{P}$. Fix a reference state $b$ and let $T = \inf\{n \geq 1 \mid X_n = Y_n = b\}$.

First, we want to show $\mathbb{P}\{T < \infty\} = 1$, i.e. the return time is finite. We begin by observing that $W_n = \{X_n, Y_n\}$ is a Markov chain on the state space $S \times S$ with transition probabilities

$$\tilde{p}(i,k)(j,l) = p(i,j)p(k,l)$$

and initial distribution

$$\mu(i,k) = \lambda(i)\pi(k).$$

Given $\mathbf{P}$ is aperiodic and irreducible, by Lemma 2.20,

$$\tilde{p}^{(n)}(i,k)(j,l) = p^{(n)}(i,j)p^{(n)}(k,l) > 0$$

for a sufficiently large $n = \max\{N(i,j), N(k,l)\}$. Thus, $\tilde{\mathbf{P}}$, the transition matrix for $W_n$, is irreducible by Definition 2.12.

We observe $\tilde{\mathbf{P}}$ also has stationary distribution

$$\tilde{\pi}(i,k) = \pi(i)\pi(k).$$

Thus $\tilde{\mathbf{P}}$ is positive recurrent by Theorem 2.18. Since $\tilde{\mathbf{P}}$ is irreducible and positive recurrent, it follows by Definition 2.16, $\mathbb{P}\{T < \infty\} = 1$.

Now, we want to show $\mathbb{P}\{X_n = i\}$ converges to $\pi(i)$ as $n$ approaches $\infty$.

First, we begin by demonstrating

$$\begin{aligned}\mathbb{P}\{X_n = i\} &= \mathbb{P}\{X_n = i, n \geq T\} + \mathbb{P}\{X_n = i, n < T\}\\ &= \mathbb{P}\{Y_n = i, n \geq T\} + \mathbb{P}\{X_n = i, n < T\},\end{aligned}$$

since $X_T = Y_T$ by construction and given both transition kernels for $X_n$ and $Y_n$ are the same.

Then, $\mathbb{P}\{X_n = i\}$

$$= \mathbb{P}\{Y_n = i \mid n \geq T\} + \mathbb{P}\{X_n = i \mid n < T\}$$
$$= \mathbb{P}\{Y_n = i \mid n \geq T\} + \mathbb{P}\{Y_n = i \mid n \geq T\} - \mathbb{P}\{Y_n = i \mid n \geq T\} + \mathbb{P}\{X_n = i \mid n < T\}$$
$$= \mathbb{P}\{Y_n = i\} - \mathbb{P}\{Y_n = i \mid n \geq T\} + \mathbb{P}\{X_n = i \mid n < T\}$$
$$= \pi(i) - \mathbb{P}\{Y_n = i, n \geq T\} + \mathbb{P}\{X_n = i \mid n < T\},$$

given the intial distribution $Y_0$ is the invariant distribution $\pi$.

Now, $\mathbb{P}\{Y_n = i\} \geq \mathbb{P}\{Y_n = i \mid n < T\}$. Furthermore, as $n$ approaches $\infty$, $\mathbb{P}\{n < T\}$ converges to $\mathbb{P}\{T = \infty\} = 0$ since $\mathbb{P}\{T < \infty\} = 1$. Since

$$0 \geq \mathbb{P}\{Y_n = i\} \geq \mathbb{P}\{Y_n = i \mid n < T\} = 0,$$

we can conclude $\mathbb{P}\{Y_n = i\}$ converges to 0 as $n$ approaches $\infty$. Similarly, $\mathbb{P}\{X_n = i\}$ must converge to 0, and so it becomes evident that as $n$ approaches $\infty$,

$$\mathbb{P}\{X_n = i\} = \pi(i) - \mathbb{P}\{Y_n = i \mid n \geq T\} + \mathbb{P}\{X_n = i \mid n < T\} = \pi(i)$$

from our earlier equation. □

**Remark 2.22.** The argument above establishes convergence via a coupling argument from Vincent Doblin (1937). In the special case of finite state spaces, an alternative proof uses the Perron–Frobenius Theorem from linear algebra, providing another route to show that $\mathbf{P}^n$ converges entry-wise to $\pi$.

## 3. Time Reversal, Detailed Balance, and Reversibility

We now develop the notion of reversibility, a property that will underlie the construction of Markov Chain Monte Carlo algorithms.

**Definition 3.1** (Detailed Balance)**.** A distribution $\pi$ on states and transition kernel $\mathbf{P}$ satisfy **detailed balance** if

$$\pi(i)\, p(i,j) \;=\; \pi(j)\, p(j,i),$$

for all $i, j \in S$.

**Definition 3.2** (Reversibility)**.** A Markov chain with transition kernel $\mathbf{P}$ and stationary distribution $\pi$ is said to be **reversible** with respect to $\pi$ if

$$\pi(i)\, p(i,j) = \pi(j)\, p(j,i) \quad \text{for all } i, j \in S.$$

That is, the chain and $\pi$ satisfy the detailed balance condition.

To offer some intuition, for a reversible chain, if you start from its stationary state (the stable long-term distribution), watching a recording of the chain's movements would look equally plausible forwards and backwards. In other words, the chain's dynamics are indistinguishable under time reversal, and hence the chain is "reversible".

Equivalently, another way to view this is as the flow of probability between any two states is perfectly balanced: in the long run, the expected traffic from $i$ to $j$ matches exactly the traffic from $j$ back to $i$. This perfect balance across every pair of states is what makes the forward and backward pictures of the chain look similar.

**Proposition 3.3.** *If $\pi$ satisfies detailed balance with respect to $\mathbf{P}$, then $\pi$ is a stationary distribution for $\mathbf{P}$.*

*Proof.* Suppose $\pi$ satisfies detailed balance. Fix $i \in S$. For any $j \in S$,

$$\sum_{i\in S} \pi(i)p(i,j) \;=\; \sum_{i\in S} \pi(j)p(j,i) \;=\; \pi(j)\sum_{i\in S} p(j,i).$$

Since $\mathbf{P}$ is a transition kernel, $\sum_{i\in S} \pi(i)p(i,j) = \pi(j)$. Thus, $\pi = \pi\mathbf{P}$, and so $\pi$ is stationary by Definition 2.9. □

In practice, reversibility is usually verified by checking the detailed balance condition. This amounts to confirming that the probability flow between each pair of states is perfectly balanced. For MCMC, reversibility is fundamental: if we can show that the target distribution $\pi$ satisfies detailed balance with respect to the transition kernel we have constructed, then $\pi$ is automatically stationary. In other words, detailed balance provides a direct way to guarantee that the Markov chain we build has the desired sampling distribution as its long-run equilibrium.

**Remark 3.4.** Not every Markov chain is reversible, but reversibility is often imposed in the design of MCMC algorithms since the detailed balance condition provides a tractable way to verify stationarity.

To build some intuition, let's consider two examples. First, a simple random walk on an undirected graph is reversible (see Example 5, Chapter 1 of [4]): if you start in equilibrium, watching the walk backward in time looks the same as watching it forward, since steps along edges have the same chance in either direction.

In contrast, imagine a chain that moves deterministically around a three–state cycle $1 \to 2 \to 3 \to 1 \to \cdots$. This chain is not reversible: if you filmed it and played the video backward, the motion would clearly look different, always moving in the opposite direction around the cycle.

These two cases highlight the meaning of reversibility: in a reversible chain, forward and backward dynamics are indistinguishable when viewed in equilibrium.

## 4. MCMC: The Metropolis-Hastings Algorithm

We now turn to Markov Chain Monte Carlo (MCMC), a powerful application of Markov chain theory.

To gain some fundamental insight into the use of Markov Chain Monte Carlo, we will first begin by considering a hypothetical problem.

**Example 4.1** (Island Problem)**.** Suppose you want to understand how a population is spread across a group of islands. Each island represents a possible state of the system, and the "true" long-term distribution of people across the islands is what you want to know—that is your target distribution $\pi$. The difficulty is that you cannot directly measure how many people live on each island. Instead, the only option is to design rules for how individuals move between islands.

One approach to determining the population distribution is to invent a Markov chain: rules that specify how a person moves from one island to another. For example, a traveler might propose to go to a random neighboring island. Sometimes the move is accepted, other times the traveler stays put, depending on how populous the new island is. By carefully crafting these acceptance rules, you can make sure that, in the long run, the flow of people from one island A to another island B is exactly balanced by the flow from B back to A. This condition, detailed balance, ensures that the overall pattern of movement settles into the correct equilibrium.

Once the rules are in place, you let the traveler wander. At first, the traveler's steps still reflect where they began rather than the true balance of people across the islands, so this initial stretch of the journey is treated as the burn-in period and the observations are discarded. After enough time, the traveler's location at a random step is approximately distributed according to $\pi$. However, if you check their location at every single step, the samples are highly correlated/dependent — watching where they are now tells you a lot about where they will be next. To get samples that behave more like independent snapshots, you either restart the journey with a short burn-in before each recorded observation, or more commonly, you let the traveler keep walking but only record their location every so often. In this way, each saved stop along the journey gives you an (approximately) independent glimpse into the true population distribution.

In this way, Markov Chain Monte Carlo can be understood as simulating the journey of a wandering traveler. Rather than jumping directly to the answer, you design fair movement rules so that the traveler naturally spends time in each state in the right proportions. Over time, the path they trace out reveals the distribution you were trying to uncover.

In general, the idea of Monte Carlo methods is to use random sampling to approximate complicated quantities: for example, producing samples from a distribution in order to estimate its mean or variance. Classical Monte Carlo assumes we can generate independent samples directly. In many important problems, however, the target distribution is too complex or high-dimensional to permit such direct sampling. MCMC extends the Monte Carlo philosophy by instead constructing a Markov chain whose long-run equilibrium distribution is exactly the distribution we wish to sample from. Running the chain then yields approximate samples from the target law.

A classical motivating example comes from statistical physics and is known as the *Ising model*. This is a simplified model of ferromagnetism: certain metals can be modeled as lattices of atoms, each bearing a "spin" that may point up or down. Different spin configurations correspond to different total energies of the system, determined by whether neighboring spins align or oppose each other. Physical principles, formalized by the Boltzmann distribution, dictate that low-energy configurations are more likely to occur, while high-energy ones are suppressed. Mathematically, the probability of a configuration $x$ is proportional to $\exp(-H(x))$, where $H(x)$ is the Hamiltonian (energy) of $x$. The difficulty is that the number of possible spin configurations increases exponentially with the size of the lattice. Even though the probability of a single configuration can be computed from the Hamiltonian, the full probability distribution over all $2^{|\text{lattice}|}$ possibilities is intractable to describe or sample from directly, hence why we turn to MCMC methods. Namely, the Metropolis algorithm (Metropolis et al., 1953) was introduced precisely to address this issue: it constructs a Markov chain whose long-run behavior reflects the probability distribution, so that representative configurations can be generated by simulating the chain. Hastings (1970) later generalized the method to allow for more flexible proposal mechanisms, leading to the modern Metropolis–Hastings algorithm.

The idea behind this algorithm is simple: we cannot sample directly from the complicated distribution $\pi$, but we can design local moves that explore its state space correctly in the long run.

**Definition 4.2** (Metropolis–Hastings Algorithm)**.** Let $\pi$ be a target distribution on a finite state space $S$. Suppose the states are the vertices of a connected, undirected graph $G$ with maximum degree $r$. For each pair of nodes $i, j \in S$, let $g(i,j)$ denote the probability of proposing a move from $i$ to $j$. A common default choice is:

$$g(i,j) = \begin{cases} \frac{1}{r}, & \text{if } i \text{ and } j \text{ are adjacent,} \\ 0, & \text{if } i \text{ and } j \text{ are not adjacent and } i \neq j, \\ 1 - \frac{\deg(i)}{r}, & \text{if } i = j. \end{cases}$$

This ensures that the total outgoing probability from each node sums to one. (Other choices of $g(i,j)$ are possible; this specification simply provides one convenient example.)

The Metropolis–Hastings algorithm defines a Markov chain $X_t$ as follows:

(1) **Initialization.** Pick an initial state $X_0 \in S$ and set $t = 0$.
(2) **Iteration.** Given the current state $X_t = i$:
   (a) *Proposal.* Generate a candidate state $j$ according to $g(i,j)$.
   (b) *Acceptance.* Compute

$$A(i,j) \;=\; \min\!\left(1,\; \frac{\pi(j)}{\pi(i)} \cdot \frac{g(j,i)}{g(i,j)}\right).$$

   Generate a random number $u$ from $[0,1]$, chosen uniformly.
   - If $u \le A(i,j)$, accept the proposal and set $X_{t+1} = j$.
   - If $u > A(i,j)$, reject the proposal and set $X_{t+1} = i$.

(3) **Increment.** Increase $t$ to $t+1$ and repeat.

Now, we observe the transition probabilities of the chain are

$$p(i,j) = g(i,j)\,A(i,j), \quad j \neq i, \qquad p(i,i) = 1 - \sum_{j \neq i} p(i,j).$$

Thus, we can verify the Metropolis–Hastings algorithm indeed defines a Markov chain, since the distribution of the next state depends only on the current state and not on the past history. At each step, a candidate state $j$ is proposed according to $g(i,j)$, which depends only on the present state $i$, and is then accepted or rejected using a rule based solely on $i$ and $j$. The algorithm yields well-defined transition probabilities from $i$ to any $j$, so the process satisfies the Markov property.

We now prove that the Metropolis–Hastings algorithm indeed produces a chain with stationary distribution $\pi$. The essential idea is that the acceptance probability is chosen precisely so that detailed balance holds.

**Proposition 4.3** (Detailed Balance for Metropolis–Hastings)**.** *The distribution $\pi$ satisfies detailed balance with respect to the Metropolis–Hastings kernel* $\mathbf{P}$.

*Proof.* Fix $i, j \in S$ with $i \neq j$. By definition of $\mathbf{P}$,

$$\mathbf{P}(i,j) \;=\; g(i,j)\,\min\!\left(1,\; \frac{\pi(j)}{\pi(i)} \cdot \frac{g(j,i)}{g(i,j)}\right).$$

Hence

$$\pi(i)\mathbf{P}(i,j) = \pi(i)g(i,j)\,\min\!\left(1,\; \frac{\pi(j)g(j,i)}{\pi(i)g(i,j)}\right).$$

If $\pi(j)g(j,i) \leq \pi(i)g(i,j)$, then the minimum equals the ratio, so

$$\pi(i)\mathbf{P}(i,j) = \pi(i)g(i,j)\,\frac{\pi(j)g(j,i)}{\pi(i)g(i,j)} = \pi(j)g(j,i) = \pi(j)P(j,i).$$

If instead $\pi(j)g(j,i) > \pi(i)g(i,j)$, then the minimum is 1, so

$$\pi(i)\mathbf{P}(i,j) = \pi(i)g(i,j) = \pi(j)g(j,i)\,\frac{\pi(i)g(i,j)}{\pi(j)g(j,i)} = \pi(j)\mathbf{P}(j,i).$$

In either case $\pi(i)P(i,j) = \pi(j)P(j,i)$, so detailed balance holds. □

**Corollary 4.4.** *The distribution $\pi$ is stationary for the Metropolis–Hastings chain.*

*Proof.* The claim follows immediately from detailed balance (Proposition 4.3) and the fact that any reversible distribution is stationary (Proposition 3.3). □

**Proposition 4.5.** *The stationary distribution $\pi$ for the Metropolis-Hastings chain is unique.*

*Proof.* We note that the Metropolis–Hastings chain is irreducible under the standing assumption that the proposal kernel is chosen so that every state is accessible from every other, i.e. the underlying proposal graph is connected.

Moreover, by Definition 2.15, the chain is aperiodic since events such as $\{X_{t+1} = X_t\}$ occur with positive probability due to the construction of the acceptance–rejection step. In other words, for some state $i$, $p(i,i) > 0$, which implies the period of $i$ is $d(i) = 1$ by Definition 2.13. By Proposition 2.14 this suffices to conclude that the period of the chain is one, and therefore the Metropolis-Hastings chain is aperiodic.

Thus, the Metropolis-Hastings chain is both irreducible and aperiodic, and therefore ergodic by Definition 2.19, given the Metropolis-Hastings chain is finite. By Theorem 2.21 it follows that the stationary distribution $\pi$ is unique, and moreover the probability distribution of $X_n$ converges to $\pi$ as $n$ approaches $\infty$, regardless of the starting state. □

This argument illustrates the essential mechanism behind the Metropolis–Hastings algorithm. The acceptance rule enforces reversibility with respect to $\pi$, so that detailed balance holds and $\pi$ is stationary. Combined with the fact that the chain is ergodic, Theorem 2.21 ensures that the distribution of the chain converges uniquely to $\pi$ regardless of the starting state. Thus, by running the chain we obtain samples that asymptotically follow the desired distribution. This captures the central idea of MCMC: by imposing simple local balance conditions, one can design Markov chains whose long-run behavior faithfully reproduces complex global distributions that would be otherwise intractable to sample from directly.

## 5. Conclusion

In this paper, we developed the theory of discrete-time Markov chains and applied it to the construction of Markov Chain Monte Carlo methods, with emphasis on the Metropolis–Hastings algorithm. Beginning with definitions and core properties such as irreducibility, recurrence, and the existence of stationary distributions, we established the ergodic theorem for Markov chains as a foundation for understanding convergence. Building on these principles, we showed how the acceptance-rejection step of Metropolis–Hastings enforces detailed balance, guaranteeing that the target distribution is stationary for the chain. This demonstrates the central idea of MCMC: by enforcing local balance conditions, we can generate global samples from distributions that are otherwise intractable to simulate directly.

Looking forward, a natural next step is to study the efficiency of MCMC methods. Although we have established convergence in principle, in practice one must ask how long it takes for the chain to become well-approximated by its stationary distribution. This "mixing time" depends intricately on the structure of the chain and can often be analyzed through spectral properties of the transition matrix. Understanding and bounding mixing times remains a central challenge in both the theory and application of MCMC, with implications for statistical physics, Bayesian inference, and randomized algorithms more broadly.

## Acknowledgments

I would like to give an enormous thank you to my mentor, Jake Zummo. He was incredibly helpful in guiding me and in answering my many questions, and our meetings provided me many valuable insights. I would not have been able to complete this paper without his support. As a result of our conversations, my interest in probability theory and its applications has grown immensely. I would also like to thank Professor May for organizing a terrific REU and for carefully reading our papers, and I am grateful to Professor Rudenko and Professor Babai for their engaging and insightful lectures. Lastly, I would like to thank the other participants for helping create such an extraordinary experience.

## References

[1] Avrim Blum, John Hopcroft, and Ravindran Kannan. Foundations of Data Science. Cambridge University Press. 2020. http://ttic.edu/blum/book.pdf

[2] Christian P. Robert and George Casella. Monte Carlo Statistical Methods, Second Edition. Springer. 2004.

[3] Geofrey Grimmett and David Stirzaker. Probability and Random Processes, Third Edition. Oxford University Press. 2001.

[4] Gregory F. Lawler. Introduction to Stochastic Processes, Second Edition. Chapman & Hall. 2006.

[5] Gregory F. Lawler. Random Explorations. American Mathematical Society Student Mathematical Library. Volume 98. 2022.

[6] James Norris. Markov Chains. https://www.statslab.cam.ac.uk/~rrw1/markov/M.pdf

[7] Persi Diaconis. The Markov Chain Monte Carlo Revolution. https://math.uchicago.edu/~shmuel/Network-course-readings/MCMCRev.pdf

[8] Steven Lalley. Markov Chains: Basic Theory. STAT 383: Measure-Theoretic Probability 2. University of Chicago, Spring 2018. https://galton.uchicago.edu/~lalley/Courses/383/MarkovChains.pdf

[9] Steven Lalley. Introduction to Markov Chain Monte Carlo. STAT 313: Stochastic Processes II. University of Chicago, Spring 2013. https://galton.uchicago.edu/~lalley/Courses/313/ProppWilson.pdf