Indian Institute of Technology Bombay
Department of Electrical Engineering

Handout 20 EE 325 Probability and Random Processes
Lecture Notes 13 March 24, 2019

1 Markov Chains

Till now, we have been dealing mostly with random variables and random vectors. In between we also encountered some questions on repeated trials of an experiment. In one of the simple examples, we repeatedly tossed a fair coin and found various probabilities under the assumption that the trials were IID. Such time-indexed sequences of random variables belong to the class of general **discrete-time stochastic processes** (DTSP), or *random processes*. However, they are not the only ones in the DTSP class. Let us now delve a bit more deeper into the theory of stochastic processes, and equip ourselves to deal with sequences which are not simply IID.

Recall the IID coin toss model, where HEAD occurs with probability p. This repeated trials can be pictorially modeled by a state-machine as shown below.

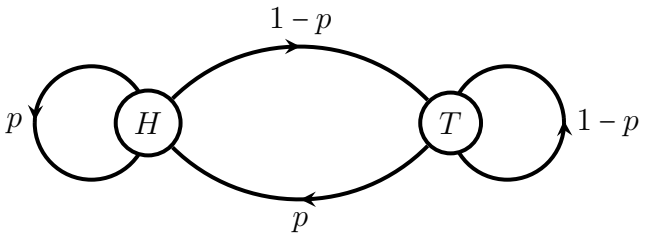

Figure 1: State diagram for repeated tosses of a coin

In figure, no matter what your current state is or past states were, the next toss takes you to the state H with probability p.

While IID sequences are definitely a very useful class of sequences which can model several processes, they are grossly insufficient in the real world. In IID trials, the probability law governing the RV stays identically the same throughout, and any finite number of joint events are governed by the product of their probabilities. This is clearly visualized in Figure 1, where the performed actions at each state does not depend on the state. Real world signals which are of interest to us have more predictable characteristics. For example, consider a temperature reading made outside our building every hour using a digital meter. Certainly this is not an iid process, as there is a lot of dependence between two successive readings. More precisely, given the current reading, the probability law governing the temperature random variable of the next reading can be different from the case where we have no access to the current reading. Some other examples of non-IID processes are

- the charge across a capacitor at integer instants.
- the samples at the output of a digital filter.
- the number of customers waiting at a ticket reservation counter at every hour.

The theory that we are about to learn here is not a panacea to the general version of all these problems. Still, it encompasses a large class of useful models, making it extremely important from an engineering perspective. We will unravel the theory behind models which admit some form of a state-variable representation as in Figure 1, however the performed action at each state can depend on the current state. So the graph may not look the same from different states, and there may not be any limit on the number of states, i.e. we can consider a countable state-space. Let us first start with the definition of a discrete-time stochastic process (DTSP). We consider discrete random variables throughout this chapter, and will not explicitly state it every-time. The theory, however, can be extended to general RVs.

Definition 1 *A sequence of RVs $X_n, n \geq 0$ is a DTSP on a countable state-space E, if each X_n is a random variable in E.*

Each possible sequence is called a trajectory or a sample-path of the processes. Thus a DTSP can be visualized as the choice of a random path, i.e. each ω correspond to a trajectory from the *cylinder set* of all possible sample-paths. A DTSP is statistically characterized by its finite dimensional distributions, i.e. we need to specify the probability $P(\bigcap_{i\in I} X_i = x_i)$ for all finite index-sets $I \subset \mathbb{N}$. It is usual to refer to the indices in the subscript of the process as time-instants, or steps, or transitions of the process.

Markov Chains are a particular class of DTSPs which have wide applications in practice. The beauty of these sequences is literally captured in the expression 'conditional independence of past and future given the present'.

Definition 2 *A DTSP $X_n, n \geq 0$ on the state-space E is called a Markov Chain if*

$$P(X_{k+1} = j | X_0 = i_0, X_1 = i_1, \cdots, X_k = i_k) = P(X_{k+1} = j | X_k = i_k),\ \forall k. \tag{1}$$

In other words, the probability law for X_{k+1} depends only on the current state X_k, when the latter is given. Notice that X_{k+1} and X_{k-1} are not always independent for a Markov chain, they are conditionally independent given X_k.

Definition 3 *A Markov Chain on E is called homogeneous if*

$$P(X_{k+1} = j | X_k = i) = P(X_{m+1} = j | X_m = i),\ \forall k, m \in \mathbb{N}. \tag{2}$$

We will use the abbreviation HMC for a homogeneous Markov Chain. A HMC can be specified completely by the probability law of X_0, along with the values

$$p_{ij} = P(X_1 = j | X_0 = i),\ \forall i, j \in E. \tag{3}$$

To see this, apply Baye's rule to expand the finite dimensional distributions as,

$$\begin{aligned} P(X_{i_0} = j_0, \cdots, X_{i_k} = j_k) &= P(X_{i_0} = j_0, X_{i_1} = j_1) \prod_{l\geq 2} P(X_{i_l} = j_l | X_{i_0} = j_0, \cdots, X_{i_{l-1}} = j_{l-1}) \\ &= P(X_{i_0} = j_0) P(X_{i_1} = j_1 | X_{i_0} = j_0) \prod_{l\geq 2} P(X_{i_l} = j_l | X_{i_{l-1}} = j_{l-1}) \\ &= P(X_{i_0} = j_0) \prod_{l\geq 1} P(X_{i_l} = j_l | X_{i_{l-1}} = j_{l-1}) \\ &= P(X_{i_0} = j_0) \prod_{l\geq 1} p_{j_{l-1} j_l} \end{aligned} \tag{4}$$

Thus we need to specify only the initial distribution and the probability of successive transitions, given by (3). The entries p_{ij} are called the **transition probability** from

state i to state j, and the matrix $\mathbb{P}$ having entries as $(\mathbb{P})_{i,j} = p_{ij}$ is known as the *transition probability matrix*. The matrix representation allows the elegant representation of various relations. To facilitate this, let $\bar{\nu}_n$ represent the array with

$$\bar{\nu}_n(i) = P(X_n = i), i \in E.$$

By convention, we will consider arrays along the column. Thus, when the array dimension is finite, $\bar{\nu}_n$ is nothing but a column-vector. Let $\bar{\nu}_n^T$ denote the transpose of $\bar{\nu}_n$, which in fact is a row-array. We can then write,

$$\begin{aligned}
\bar{\nu}_{n+1}(j) &= P(X_{n+1} = j) && (5)\\
&= \sum_{i\in E} P(X_n = i)P(X_{n+1} = j|X_n = i) && (6)\\
&= \sum_{i\in E} \bar{\nu}_n(i)p_{ij}. && (7)
\end{aligned}$$

In the matrix format,

$$\begin{aligned}
\bar{\nu}_{n+1}^T &= \bar{\nu}_n^T \mathbb{P} && (8)\\
&= \bar{\nu}_{n-1}^T \mathbb{P} \times \mathbb{P} && (9)\\
&= \bar{\nu}_0^T \mathbb{P}^{n+1}, && (10)
\end{aligned}$$

which re-asserts what we mentioned earlier, that the HMC is specified by an initial distribution $\bar{\nu}_0$ and the transition matrix $\mathbb{P}$.

A ubiquitous first example for a Markov Chain is the so called random walk on $\mathcal{Z}$.

Example 1 *Consider a sequence of steps $X_n, n \geq 0$ taken by a random walk on $\mathbb{Z}$, where X_n is incremented or decremented by one unit, based on independent tosses of a coin with $P(HEAD) = p$. Show that $X_n, n \geq 0$ is a Markov Chain. Is this a HMC?*

Solution: It is easy to verify that

$$P(X_{n+1} = j|X_0 = i_0, \cdots, X_n = i_n) = p\mathbb{1}_{\{j=i_n+1\}} + (1-p)\mathbb{1}_{\{j=i_n-1\}} = P(X_{n+1} = j|X_n = i_n). \quad (11)$$

Clearly, it is homogeneous since the transition probability from i to j stays the same, irrespective of when you reach state i.

■

Evaluating the conditional probability of a DTSP to check for markovity every time can be a mundane task. To circumvent this, there are two recurrence relations which can make life easier for us. These relations only test for HMCs, the converse is not true, i.e. all HMCs may not admit such relations that we prescribe.

Theorem 1 *Recurrence 1: Consider the recurrence relation*

$$X_{n+1} = f(X_n, Z_{n+1}),$$

where $Z_n, n \geq 0$ is IID independent of X_0. Then $X_n, n \geq 0$ is a HMC.

Proof

$$\begin{aligned}
&P(f(X_n, Z_{n+1}) = j|X_n = i, X_{n-1} = i_{n-1}, \cdots, X_0 = i_0)\\
&\qquad = P(f(i, Z_{n+1}) = j|X_n = i, X_{n-1} = i_{n-1}, \cdots, X_0 = i_0) && (12)\\
&\qquad = P(f(i, Z_{n+1}) = j)\\
&\qquad = P(P(f(X_n, Z_{n+1}) = j)|X_n = i),
\end{aligned}$$

where the second equality states that the probability in (12) is determined solely by the law of Z_{n+1}, which is independent of $X_k, 0 \le k \le n$.

■

It is now easy to apply this result to the random walk in Example 1. Here the HMC obeys

$$X_{n+1} = X_n + Z_n,$$

where $Z_n \in \{-1, +1\}$ with $P(Z_n = 1) = p$. There are several other examples where this recurrence relation readily applies.

Our second recurrence relation is more general and a very useful tool.

Theorem 2 *Recurrence 2: Consider a collection of probability laws $p_i(x), x \in E$, one law for each $i \in E$, along with the relation*

$$X_{n+1} = f(X_n, Z_{n+1}),$$

where Z_{n+1} is independent of $X_{n-1}, \cdots, X_0, Z_n, \cdots, Z_1$ given X_n. Furthermore if $Z_{n+1} \sim p_i(\cdot)$ whenever $X_n = i$, then $\{X_n\}, n \ge 0$ is a HMC with

$$p_{ij} = P(f(i, Z_1) = j | X_0 = i).$$

Proof: Similar to the previous theorem.

■

Note: I have changed the definition slightly from the one given in class, somehow I thought my statement implied the same conditional law, but it has to be explicitly stated, apologies and please correct the notes.

The generality of the latter theorem comes from the fact that we only demand the conditional independence Z_{n+1} from the past given X_n. The following example shows the utility of this theorem.

Example 2 ***EhrenFest's Urn Model****: N balls are distributed among 2 boxes, say box A and box B. At any step, a ball is uniformly chosen. This ball is then moved from the incumbent box to the other. The experiment is repeated at every instant. Show that the number of balls in Box A is a HMC and sketch its state-diagram.*

Solution: Let $X_n, n \ge 0$ denote the number of balls in Box A at step n. The state-diagram is shown below. The state-space graph itself will guarantee us that the underlying process

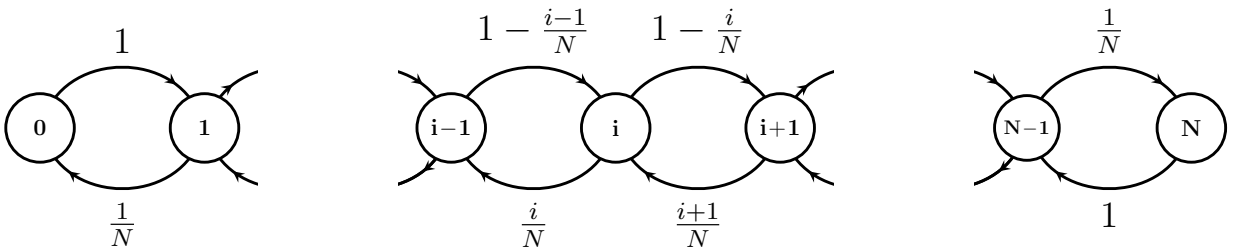

Figure 2: EhrenFest's Urn State-space

is a HMC, as the transition probabilities from any given state depends only on the state

and are time-invariant. While our first recurrence relation does not fit this process well, the second one naturally does. To see this, observe that

$$X_{n+1} = X_n + Z_{n+1},$$

where given $X_n = i$, the random variable Z_n has the law

$$Z_{n+1} = \begin{cases} +1 \text{ w.p } 1 - \frac{i}{N} \\ -1 \text{ w.p } \frac{i}{N} \end{cases}$$

Thus $X_n, n \geq 0$ is a HMC by our second recurrence relation.

■

1.1 Irreducible Markov Chains

Irreducibility is a graph-theoretic property, which physically signifies the *connectivity* of a graph. For Markov chains, irreducibility of the state-space is a slightly more general notion, as the state-space can be countably many.

Definition 4 *A HMC $X_n, n \geq 0$ on a state-space E is irreducible, if there exists a trajectory which contains states i and j, for every pair $(i, j) \in (E \times E)$.*

The definition implies that every state can be reached from every other state through a sequence of steps of the chain. In terms of the transition probability, we can state that for some n

$$(\mathbb{P}^n)_{i,j} > 0.$$

Thus, there exists a number n such that there is a positive probability of reaching state j from state i in n steps. We will use the notation

$$p_{ij}(n) \stackrel{\triangle}{=} P(X_n = j | X_0 = i), \tag{13}$$

to denote the probability of reaching state j in n steps starting from the initial state i, which is nothing but $(\mathbb{P}^n)_{i,j}$. This is also known as the $n-$ step transition probability.

Using the $n-$step transition probabilities,

$$\bar{\nu}_n(j) = \sum_i p_{ij}(n)\nu_0(i). \tag{14}$$

What will happen to $\bar{\nu}_n$ as the time progresses or n gets large. Contrary to what we may expect, this question has a natural answer where $\bar{\nu}_n$ becomes independent of $\bar{\nu}_0$ in the limit of n. This leads to the notion of equilibrium or steady state which we explain in the next section.

1.2 Stationary Distribution

The stationary distribution plays an important role in the theory of Markov chains. This is also known as equilibrium distribution or steady-state distribution. It is natural to expect that once you are in steady state, you stay in steady state, forever. We will define Π with this property.

Definition 5 *For a HMC with transition matrix* $\mathbb{P}$*, if*

$$\Pi^T = \Pi^T \mathbb{P}, \tag{15}$$

for some non-negative vector which sums to one, then Π *is called a stationary distribution of* $\mathbb{P}$*. The equation in* (15) *is known as the* ***global balance*** *equation.*

For the time being, let us assume that a stationary distribution exists, the existence and uniqueness will be discussed in later sections. Just like a chemical reaction inching close to equilibrium as time goes on, a Markov chain also will get to a stationary distribution, when run for a long enough time. The time taken to reach steady-state is known as the *burn-in* time. It turns out that the steady state distribution only depends on the transition matrix, and not on the initial distribution. We do not prove this statement, the usual proofs in literature employ elegant coupling arguments of multiple HMCs (see P. Bremaud, Markov Chains). Thus, no matter what the initial distribution is, in steady state the HMC will have a distribution solely determined by $\mathbb{P}$. We denote the stationary distribution by Π.

Observe that we did abuse the language in using 'the' stationary distribution, since it is not clear apriori whether there are multiple stationary distributions. Indeed there are many HMCs with multiple stationary distributions. Nevertheless, the important class of 'irreducible HMCs' have a unique stationary distribution whenever it exists(see P. Bremaud, Markov Chains). We will also learn more on the physical meaning of stationary distribution by relating it to the so called invariants of a chain. Before we progress to these aspects, it is instructive to compute the stationary distribution for some example chains to gain more insight. The computations typically involve solving the global balance equations, which can be enumerated as

$$\Pi(j) = \sum_{i \in E} \Pi(i) p_{ij} \tag{16}$$

Notice that (15) is not a set of independent equations. Indeed for the vector $\bar{1}$ will all ones,

$$\Pi^T \bar{1} = \Pi^T \mathbb{P} \times \bar{1} = \Pi^T \bar{1}, \tag{17}$$

since $\mathbb{P}$ is a stochastic matrix, i.e. the entries are non-negative and each row adds to unity. Clearly one of the global balance equations is redundant. However, we can use the fact that $\sum_{i \in E} \Pi(i) = 1$, to solve Π unambiguously for irreducible HMCs. Keep this in mind while solving the set of equations.

Example 3 *Find the stationary distribution of the two-state chain as shown below.*

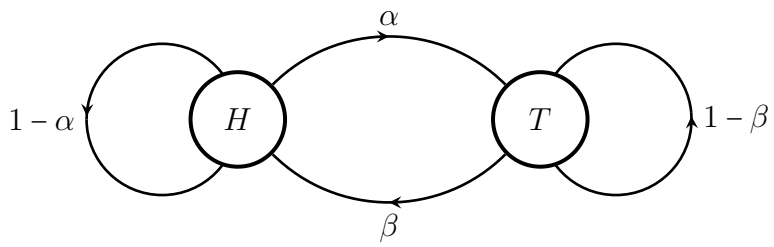

Figure 3: Two State HMC

Solution: Observe that the two-state chain signifies a physical process where there are two biased coins. If your current state is HEAD, a biased coin of probability α is used to determine your transition to TAILs. On the other hand, if the current state is TAILs, we will employ a biased coin of probability β for the transitions. Denoting HEADS by 0 and TAILs by 1, let us write the global balance equations.

$$\begin{aligned}\Pi(0) &= \Pi(0)(1-\alpha) + \Pi(1)\beta \\ \Pi(1) &= \Pi(0)\alpha + \Pi(1)(1-\beta).\end{aligned}$$

Clearly, summing the two equations will give us a redundant equation. However, substituting $\Pi(0) + \Pi(1) = 1$ in the first equation will imply

$$\Pi(0) = \Pi(0)(1-\alpha) + (1-\Pi(0))\beta.$$

Thus,

$$\Pi(0) = \frac{\beta}{\alpha+\beta} \text{ and } \Pi(1) = \frac{\beta}{\alpha+\beta}.$$

An alternate method is to solve for the left eigen-vectors of the matrix

$$\mathbb{P} = \begin{bmatrix} 1-\alpha & \alpha \\ \beta & 1-\beta \end{bmatrix}, \tag{18}$$

corresponding to the eigen-value 1. Remember to compute the left eigen-vector. Right eigen-vectors may not correspond to the stationary distribution. Right multiply $\mathbb{P}$ with the vector $\left[\frac{1}{2}\ \frac{1}{2}\right]$ and see for yourself.

■

Example 4 *Compute the stationary distribution for the Ehrenfest's urn model.*

Solution: This shows another useful technique to solve the global balance equations. In here, we start with one state, say state 0 and express $\Pi(i), i \neq 0$ in terms of $\Pi(0)$. In the end, choose $\Pi(0)$ such that Π is a probability distribution. Let us start by writing

$$\Pi(0) = \Pi(1)\frac{1}{N}.$$

Thus,

$$\Pi(1) = N\Pi(0).$$

Now

$$\Pi(1) = \Pi(0) + \Pi(2)\frac{2}{N}.$$

From this

$$\begin{aligned}\Pi(2) &= \frac{N}{2}\left(\Pi(1) - \Pi(0)\right) \\ &= \frac{N(N-1)}{1\times 2}\Pi(0) \\ &= \binom{N}{2}\Pi(0).\end{aligned}$$

Continuing this recursion

$$\Pi(k) = \binom{N}{k} \Pi(0), 1 \leq k \leq N.$$

Since Π is a probability distribution,

$$\begin{aligned} 1 &= \sum_{i=1}^{N} \Pi(i) \\ &= \sum_{i=1}^{N} \binom{N}{i} \Pi(0) \\ &= 2^N \Pi(0). \end{aligned}$$

Finally,

$$\Pi(i) = \frac{\binom{N}{i}}{2^N}. \tag{19}$$

■

Let us now pursue our original goal of understanding the stationary distributions. We will relate them to a more general property of the HMC, known as 'invariants'.

2 Invariants

We will discuss mostly about irreducible HMCs in this discussion. This allows us to use the word 'the stationary distribution', whenever it exists. The term *invariant* is a generalization of the stationary distribution, i.e. an invariant x obeys

$$x^T = x^T \mathbb{P}.$$

We did not insist here on x being a probability distribution. The invariants are closely related to the return-times of a HMC to any state, and has a intuitive frequency interpretation. The return-time to state i is the number of transitions of the chain between two consecutive visits to state i. In other words, it is the length of the trajectory between two successive visits to any state.

Without loss of generality, we will take state 0 as our starting state, i.e. $X_0 = 0$ for our HMC. If you do not like this, take any other starting state and rename it as state zero. Let $T_0 = \inf\{n \geq 1 : X_n = 0\}$ denote the number of steps it takes for the HMC to return to state 0. Clearly T_0 is a function of $X_n, n \geq 0$ and a random variable. We are interested in the expected time it takes for the first return to state 0, from the same initial state. We will denote this by $\mathbb{E}_0[T_0]$, where the subscript of the expectation denotes the starting state. In a similar way, we will denote by $P_0(A)$, the conditional probability of event A, given that the HMC is initially at state 0.

All this talk about return times is to show that the stationary distribution is connected to the return times to a state. Consider a function defined for every $i \in E$ as

$$x_i = \mathbb{E}_0\Big[\sum_{n \geq 1} \mathbb{1}_{\{X_n = i\}} \mathbb{1}_{\{n \leq T_0\}}\Big]. \tag{20}$$

As mentioned earlier, the subscript 0 indicates that the initial state is $X_0 = 0$, and T_0 is the time/number-of-steps for the chain to return to zero. Thus x_i is the expected number of

times the state i is visited, in between two consecutive visits to zero. Clearly $x_0 = 1$, as the state zero is visited only while returning to it, by our definition. In order to understand (20) better, let us compute the sum $\sum_{i\in E} x_i$.

$$\begin{aligned}\sum_{i\in E} x_i &= \sum_{i\in E} \mathbb{E}_0\Big[\sum_{n\geq 1} \mathbb{1}_{\{X_n=i\}} \mathbb{1}_{\{n\leq T_0\}}\Big] \\ &= \mathbb{E}_0\Big[\sum_{n\geq 1} \mathbb{1}_{\{n\leq T_0\}} \sum_{i\in E} \mathbb{1}_{\{X_n=i\}}\Big] \\ &= \mathbb{E}_0\Big[\sum_{n\geq 1} \mathbb{1}_{\{n\leq T_0\}}\Big] \\ &= \mathbb{E}_0\big[T_0\big]. \end{aligned} \tag{21}$$

Thus $\sum_i x_i$ is the expected duration of the return-time to zero. This makes sense, since we are adding the expected hits in each state between two successive visits to zero, so the cumulant better be equal to the average length (steps) for returning to state 0.

Consider the column array $\bar{x}$ with the i^{th} entry as x_i, $i \in E$. Then we have the following theorem, which is one of the important results on Markov chains.

Theorem 3 *The $\bar{x}$ defined in (20) is an invariant to $\mathbb{P}$, i.e.*

$$\bar{x}^T = \bar{x}^T \mathbb{P}.$$

Proof: It is clear that $x_0 = 1$. For $x_i, i \neq 0$ we have, By moving the expectation operator inside the summation in (20),

$$\begin{aligned} x_i &= \sum_{n\geq 1} \mathbb{E}_0\Big[\mathbb{1}_{\{X_n=i\}} \mathbb{1}_{\{n\leq T_0\}}\Big] \\ &= \sum_{n\geq 1} P_0(X_1 \neq 0, \cdots, X_{n-1} \neq 0, X_n = i) \\ &= \sum_{n\geq 1} P_0\left(\{X_j \neq 0, 1 \leq j \leq n-1\} \bigcap \{X_n = i\}\right), \end{aligned} \tag{22}$$

where the probability on the RHS considers those outcomes where starting from state 0

1. the chain has not hit zero till $n-1$, i.e. $T_0 > n-1$.
2. at step n the chain is at state i.

Please note that in our indexing $\{j : 1 \leq j \leq 0\} = \varnothing$, i.e. we will ignore those terms where the index is zero in the above definitions, for consistency. For economy of space let us denote the RHS of (22) as,

$$\hat{p}_{0i}(n) = P_0\left(\{X_j \neq 0, 1 \leq j \leq n-1\} \bigcap \{X_n = i\}\right), i \neq 0. \tag{23}$$

i.e. starting with state 0, the probability of reaching state i in the n^{th} step, without having hit the state 0 in steps $1 \leq k \leq n-1$. Clearly, from (22),

$$x_i = \sum_{n\geq 1} \hat{p}_{0i}(n), \forall i \neq 0. \tag{24}$$

Using (20) – (22), we can write for all $i \in E$,

$$\begin{aligned} x_i &= \sum_{n\geq 1} P_0\left(\{X_j \neq 0, 1 \leq j \leq n-1\} \bigcap \{X_n = i\}\right) \\ &= P_0(X_1 = i) + \sum_{n\geq 2} P_0(X_1 \neq 0, \cdots, X_{n-1} \neq 0, X_n = i) \\ &= \mathbb{P}_{0i} + \sum_{n\geq 2} \sum_{j\neq 0} \hat{p}_{0j}(n-1)\mathbb{P}_{ji}. \end{aligned} \tag{25}$$

This last step needs some explanation. The first term on the RHS is clear, as the probability of reaching state i from state 0 in one step is indeed given by the transition probability. The second term used the 'first-step analysis' that we already learned, which is a form of Baye's rule for the underlying HMC. In particular, reaching state i at step n will imply that the chain was at some state j at time $n-1$ and the next step took it to state i. We are okay to jump to any intermediate step other than state 0. With the initial state at 0, the probability of reaching j at step $n-1$ without hitting zero in steps $\{1, 2, \cdots, n-2\}$ is indeed $p_{0j}(n-1)$, see (23). The last factor in the second term of (25) is the probability of jumping from j to i in the n^{th} step. We can now simplify things. Since $x_0 = 1$,

$$\begin{aligned} x_i &= x_0\mathbb{P}_{0i} + \sum_{n\geq 2}\sum_{j\neq 0}\hat{p}_{0j}(n-1)\mathbb{P}_{ji} \qquad (26)\\ &= x_0\mathbb{P}_{0i} + \sum_{j\neq 0}\mathbb{P}_{ji}\sum_{n\geq 2}\hat{p}_{0j}(n-1)\\ &= x_0\mathbb{P}_{0i} + \sum_{j\neq 0}\mathbb{P}_{ji}\sum_{m\geq 1}\hat{p}_{0j}(m)\\ &= x_0\mathbb{P}_{0i} + \sum_{j\neq 0}\mathbb{P}_{ji}x_j. \qquad (27) \end{aligned}$$

The last equation follows from (24). Collecting all such equations, we get

$$x^T = x^T\mathbb{P}.$$

2.1 Invariants and Stationary Distribution

Clearly the invariant of an HMC throws a lot of light on the stationary distribution Π. For an irreducible HMC, we can show that if there is a probability vector Π such that $\Pi^T = \Pi^T\mathbb{P}$, then it is unique. A proof of concept to this idea is given as an exercise in the next subsection. This will allow us to refer Π as 'the stationary distribution' of an irreducible HMC, as there is only one possible. We should add that this does not rule out the possibility of having no Π which obeys the global balance equations. In fact, the discussion below allows a quick way to deduce the cases where a stationary distribution readily exists, and this will include all the examples and models we considered so far in the lass.

Using the concept of invariants, whenever $\sum_i x_i$ is finite, we can define

$$\Pi(i) = \frac{x_i}{\sum_{i\in E} x_i}.$$

This definition ensures that

$$\Pi^T\mathbb{P} = \frac{x^T}{\sum_{i\in E} x_i}\mathbb{P} = \frac{x^T}{\sum_{i\in E} x_i} = \Pi^T.$$

Recall that we considered state 0 as the initial state. Thus

$$\Pi(0) = \frac{x_0}{\sum_i x_i} = \frac{1}{\mathbb{E}_0[T_0]}.$$

Exercise 1 *Since* 0 *is an arbitrary state, deduce that*

$$\Pi(i) = \frac{1}{E_i[T_i]},$$

where $\mathbb{E}_i[T_i]$ *is the expected time to return to state* i*, starting from state* i*.*

Let us now give a condition for the distribution Π to exist, this turns out simple. Whenever $\mathbb{E}_i[T_i]$ is finite, we can find $\Pi(i)$ for all i. For an irreducible chain, it is sufficient to check whether $\mathbb{E}_0[T_0]$ is finite.

Exercise 2 *Show that for any irreducible HMC with transition matrix $\mathbb{P}$ on a finite state-space, there exists a Π such that $\Pi^T = \Pi^T\mathbb{P}$.*

Exercise 3 *We defined earlier $\nu_n(i) = P(X_n = i)$ for an HMC $X_n, n \geq 0$. Suppose $\nu_0(i) = \Pi(i), \forall i \in E$. Find $\nu_n(i)$ for $n \geq 1$.*

So far we have worked under the assumption that a stationary distribution exists. It is now clear that whenever $\mathbb{E}_0[T_0] < \infty$, we will have a stationary distribution for the irreducible HMC. A HMC with $\mathbb{E}_0[T_0] < \infty$ is known as **positive recurrent**. Positive recurrence is synonymous with the word 'stability', and is a central notion in the theory of HMCs. All the HMCs that we considered so far belongs to the positive recurrent category. However there are HMCs which are not positive recurrent, a symmetric random walk is an example of a so called null-recurrent chain. For lack of time, we do not cover the details.

2.2 Periodicity

The stationary distribution is also commonly known as the equilibrium distribution. Let us point out a subtle difference between these two terminologies. It turns out that a probability vector Π obeying the global balance equations should be more aptly called the equilibrium distribution, as it talks about the balance of equations. Consider the two-state HMC given below.

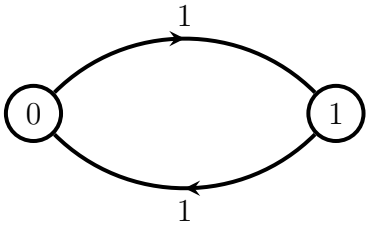

Figure 4: Cyclic Chain

Clearly this is an irreducible HMC, and we can find the equilibrium distribution $\Pi = [\frac{1}{2}, \frac{1}{2}]^T$. However, the initial state completely determines the trajectory of the chain. This happens due to the cyclic nature of the transitions. For an HMC let, $\mathcal{N}_i = \{n_1, n_2, \cdots\} \subset \mathbb{N}$ denote the collection of instants such that for each $n \in \mathcal{N}_i$, there exists a trajectory starting from state i which reaches state i at step n. The GCD of all $n \in \mathcal{N}_i$ is the period of the state i. A irreducible HMC is called aperiodic if the period of each state is 1.

Exercise 4 *Show that if $\mathbb{P}_{ii} > 0$ for some $i \in E$, then an irreducible HMC is also aperiodic.*

A well known result on Markov chains is that no matter what the initial distribution is, an aperiodic irreducible HMC will end up in a stationary distribution as n goes to infinity. In other words, for an irreducible aperiodic HMC with the stationary distribution Π, we have

$$\lim_{n\uparrow\infty} \nu_n(i) = \Pi(i).$$

This is a soft point where many students demanded a proof. Rather than working out an epsilon-delta proof by analysis, the next subsection relies on the concept of entropy in information theory. This has considerably more intuitive appeal, and also allows us to visualize the connections to the equilibrium.

2.3 Stationary Distribution: Existence and Uniqueness

We should start with a warning that this is not a complete proof of a limiting statement in the last section. Rather we will formalize some steps using the ideas borrowed from information theory, which will help you understand the uniqueness of the stationary distribution and why any initial distribution converges to a stationary distribution.

Exercise 5 *The entropy of a distribution $P(x)$ is given by*

$$H(X) = \sum_{x \in E} P(x) \log \frac{1}{P(x)}.$$

Using $\log_e(x) \leq x - 1$, show that $H(X) \geq 0$.

Exercise 6 *The relative entropy between two distributions $P(x)$ and $Q(x)$ are given by*

$$D(P||Q) = \sum_{x \in E} P(x) \log \frac{P(x)}{Q(x)},$$

where we assumed that $Q(x) = 0$ whenever $P(x) = 0$, and multiplying by zero always results in zero. Show that $D(P||Q) \geq 0$, and that equality occurs only when $P(x) = Q(x)$ for all $x \in E$.

Let us now consider an HMC with transition matrix $\mathbb{P}$. Consider two distributions $P(\cdot)$ and $Q(\cdot)$ on E. If we start the chain with initial distribution $P(\cdot)$, let $P(x_n, x_{n+1})$ be the joint distribution of states at times $(n, n+1)$. Let $\mu_n(x)$ denote $P(X_n = x)$ when we start with the initial distribution $P(x)$. Similarly, let $Q(x_n, x_{n+1})$ be the joint distribution of states at times $(n, n+1)$ if we start with the initial distribution $Q(\cdot)$. Let us denote $\nu_n(x) := P(X_n = x)$ when we start with the initial distribution $Q(x)$.

The following exercise will convince you that that the relative entropy between $\nu_n(x)$ and $\mu_n(x)$ will decrease as n goes to infinity

Exercise 7 *Show that*

$$D(P(x_n, x_{n+1})||Q(x_n, x_{n+1})) = D(\mu_n(x)||\nu_n(x)).$$

Furthermore, by another application of Baye's rule, and the fact that the relative entropy is non-negative, show that

$$D(P(x_n, x_{n+1})||Q(x_n, x_{n+1})) \geq D(\mu_{n+1}(x)||\nu_{n+1}(x)).$$

The above exercise makes it clear that

$$D(\mu_{n+1}(x)||\nu_{n+1}(x)) \leq D(\mu_n(x)||\nu_n(x)).$$

Thus if you take $\nu_0(x) = Q(x) = \Pi(x)$, for any stationary distribution Π, the relative entropy between Π and μ_n can only decrease. This will give you an idea why an HMC should eventually end up in a stationary distribution. What we have not proved is that the decrease is relative entropy is sufficiently large in each step to reach convergence as n gets large. Please look the references given for Markov Chains for further details.

3 Reversibility

In the class, we studied the global balance equations. There are also a set of conditions known as local balance equations, which helps to find the stationary distribution in an easier manner. However, this technique will only apply to the so called 'reversible' chains. Notice that not all chains are reversible though. Nevertheless, it is a good idea to try the local balance before you proceed to apply the global balance. The local balance condition for an irreducible chain is

$$\Pi(i)p_{ij} = \Pi(j)p_{ji}, \forall i, j \in E.$$

Example 5 *Show that the local balance implies the global balance equations, and thus any solution to* Π *above is the stationary distribution to the irreducible chain.*

Solution:

$$\begin{aligned}\sum_{i\in E} \Pi(i)p_{ij} &= \sum_{i\in E} \Pi(j)p_{ji} \\ &= \Pi(j) \sum_{i\in E} p_{ji} \\ &= \Pi(j),\end{aligned}$$

which is nothing but the global balance equation. Thus the local balance is a stronger condition.

Example 6 *Solve the stationary distribution for the two-state chain and also the Ehrenfest's urn model using the local balance equations.*

Solution: For the two state chain

$$\Pi(0)\alpha = \Pi(1)\beta$$

Since $\Pi(0) + \Pi(1) = 1$,

$$\Pi(0) + \Pi(0)\frac{\alpha}{\beta} = 1.$$

Thus

$$\Pi(0) = \frac{\beta}{\alpha+\beta} \text{ and } \Pi(1) = \frac{\alpha}{\alpha+\beta}.$$

Now for the urn model,

$$\Pi(i+1)\frac{i+1}{N} = \Pi(i)(1 - \frac{i}{N}).$$

Thus

$$\Pi(i+1) = \frac{N-i}{i+1}\Pi(i) = \binom{N}{i+1}\Pi(0),$$

and the result follows from this.

A physical/intuitive interpretation of reversibility is as follows. Consider an irreducible HMC with the stationary distribution Π as the initial distribution. Now, run the chain forward, let us say for n steps to obtain $X_0, \cdots, X_n$. Imagine that you repeat the same experiment and obtain another sequence $Y_0, \cdots, Y_n$. Suppose a third party takes the Y_n sequence and writes it in the reverse order to obtain $U_i, 0 \le i \le n$. Now you are shown one of the two sequences $X_1, \cdots, X_n$ or $U_1, \cdots, U_n$, and asked to identify whether it is the reversed one or not. If the chain is reversible, the best thing that you can do is to toss a fair coin, and answer without even looking at the sequence. The statistical properties for the forward and reverse runs are identical in reversible HMCs.

4 Random walk on a Graph

Consider an irreducible graph with a finite number of nodes and a given edge-set. Consider a random walk on this graph, where from each vertex, one of the neighbors is uniformly chosen. Clearly the state of this walk is a HMC, let us find the stationary distribution of this process.
Note: Though this is known as random walk on a graph, it is different from conventional random walks.

We can first check whether the local balance equations can be solved.

$$\Pi(i)\frac{1}{d_i} = \Pi(j)\frac{1}{d_j}$$

or

$$\Pi(i) = \frac{d_i}{d_j}\Pi(j)$$

Summing over i

$$\sum_{i\in E}\Pi(i) = 1 = \frac{\Pi(j)}{d_j}\sum_{i\in E} d_i.$$

Thus

$$\Pi(j) = \frac{d_j}{\sum_{i\in E} d_i},$$

gives the stationary distribution of the walk on the graph.

Exercise 8 *Random Horse in a checkerboard: Consider a horse in the chess game, which is placed at the bottom left square. Assume that the horse takes a random walk, with each position to which it can jump is taken with equal probability. Find the expected time before it gets back to the initial position.*

Exercise 9 *Consider the undirected graph with* 9 *nodes depicted below, with edges as shown by the naked links. Consider a symmetric random walk on this graph, where at each node, one of the available links is uniformly chosen to determine the next position.*

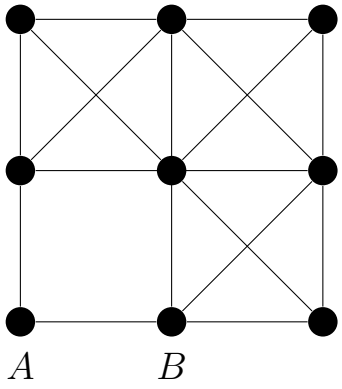

(a) *Starting from node B (see Figure), find the expected number of times node A is visited before returning to node B.*
(b) *Starting from the node B, find the expected number of steps to reach node A.*
Hint: Use the following hints only after trying your ideas, that too at your own risk. For part (a), the definition and concept of an invariant can be helpful. For part (b), a first step analysis for the expected time is perhaps useful.

5 Epilogue

The topic of Markov chains can be a course in itself. There are two courses in the department which delves much more deeper in to the theory and applications (Markov Chains and Queues, Advanced Probability). The idea of touching the topics was to not only to introduce stochastic processes, but also to illustrate the possible applications of Markov chains.

I request the reader to forgive me for not covering further details of some fundamental concepts. For example, we need much more machinery to show the convergence to equillibrium, clearly out of the scope for a first level course. This coverage is also guilty of leaving out the topic of recurrence and stability of HMCs, or the related ergodic theory. In these notes, I have attempted to address some of the questions raised in the class. We will use one of the saturdays to illustrate a particular application of the properties related to the stationary distribution, to analyse a random access protocol called ALOHA used in wireless networks.

Feel free to mail any questions that you have.