Existence and Uniqueness of the Stationary Distribution

https://asanchowk.github.io/

(This blogpost follows the exposition of Chapter 1 of [1], but in a verbose manner.)

***

Assume a *finite irreducible and aperiodic* Markov Chain $\{X_t\}_{t\geq 0}$ with transition matrix P. We prove, by construction, the existence of the stationary distribution π for every such Markov Chain.

We point the reader to this post for the notation used in this blog post. We specially point out the notation of first return time τ_z^+ for a random walk starting at $z \in \mathcal{X}$ and the element-wise formulation of the stationary distribution π for reference.

1 Setup and Overview

Let $z \in \mathcal{X}$ be any arbitrary state, and define

$$\begin{aligned}\tilde{\pi}(y) &= \mathbb{E}_z[\text{number of visits to } y \text{ before returning to } z] \\ &= \sum_{t=0}^{\infty} \Pr_z[X_t = y, \tau_z^+ > t]\end{aligned}$$

Our argument will broadly be in two steps:

1. We first show that $\tilde{\pi}P = \tilde{\pi}$.
2. $\tilde{\pi}$ is clearly not a probability distribution, so we say that $\pi = \frac{\tilde{\pi}}{\sum_{x\in\mathcal{X}} \tilde{\pi}(x)}$ is a stationary distribution on $\mathcal{X}$. We will later show that $\sum_{x\in\mathcal{X}} \tilde{\pi}(x) = \mathbb{E}_z[\tau_z^+]$, so the final stationary distribution we construct will be $\pi = \frac{\tilde{\pi}}{\mathbb{E}_z[\tau_z^+]}$.

Step 2 above shows that it is essential first to show that $\mathbb{E}_z[\tau_z^+]$ is finite. We show that in Lemma 6.2.

2 Proof of existence

Proposition 1. (Proof of existence) $\pi = \frac{\tilde{\pi}}{\mathbb{E}_z[\tau_z^+]}$ is a stationary distribution.

Proof. We begin by showing $\tilde{\pi}P = \tilde{\pi}$:

$$
\begin{aligned}
\tilde{\pi}P &= \sum_{x\in\mathcal{X}} \tilde{\pi}(x) \cdot P(x,y) \\
&= \sum_{x\in\mathcal{X}} P(x,y) \sum_{t=0}^{\infty} \Pr_z[X_t = x, \tau_z^+ > t] && \text{(from the definition of } \tilde{\pi}\text{)} \\
&= \sum_{x\in\mathcal{X}} \sum_{t=0}^{\infty} \Pr_z[X_t = x, \tau_z^+ > t] \cdot P(x,y) \\
&= \sum_{t=0}^{\infty} \sum_{x\in\mathcal{X}} \Pr_z[X_t = x, \tau_z^+ > t] \cdot P(x,y) && \text{(switching order of summations)} \\
&= \sum_{t=0}^{\infty} \sum_{x\in\mathcal{X}} \Pr_z[X_t = x, \tau_z^+ \geq t+1] \cdot P(x,y) \\
&= \sum_{t=0}^{\infty} \sum_{x\in\mathcal{X}} \Pr_z[X_t = x, X_{t+1} = y, \tau_z^+ \geq t+1] \\
&= \sum_{t=0}^{\infty} \sum_{x\in\mathcal{X}} \Pr_z[X_t = x, X_{t+1} = y, \tau_z^+ \geq t+1] \\
&= \sum_{t=0}^{\infty} \sum_{x\in\mathcal{X}} \Pr_z[X_{t+1} = y, \tau_z^+ \geq t+1] \cdot \Pr_z[X_t = x] \\
&= \sum_{t=0}^{\infty} \Pr_z[X_{t+1} = y, \tau_z^+ \geq t+1] && \left(\sum_{x\in\mathcal{X}} \Pr_z[X_t = x] = 1\right) \\
&= \sum_{t=1}^{\infty} \Pr_z[X_t = y, \tau_z^+ \geq t] \\
&= \sum_{t=1}^{\infty} \Pr_z[X_t = y, \tau_z^+ > t] + \sum_{t=1}^{\infty} \Pr_z[X_t = y, \tau_z^+ = t] \\
&= \left(\sum_{t=0}^{\infty} \Pr_z[X_t = y, \tau_z^+ > t]\right) - \underbrace{\Pr_z[X_0 = y, \tau_z^+ > 0]}_{(a)} + \underbrace{\sum_{t=1}^{\infty} \Pr_z[X_t = y, \tau_z^+ = t]}_{(b)}
\end{aligned}
$$

Recall that y is some state in $\mathcal{X}$ and in our setup $X_0 = z$. We now consider two cases:

- $y = z$. In this case the terms (a) and (b) are both equal to 1, so they cancel out.
- $y \neq z$. In this case the terms (a) and (b) are both equal to 0. Term (a) is 0 because we know $X_0 = z$ and term (b) is 0 because at time τ_z^+ we know $X_t = z$, so $\Pr_z[X_t = y, \tau_z^+ > t] = 0$.

Either case simplifies to

$$\sum_{t=0}^{\infty} \Pr_z[X_t = y, \tau_z^+ > t] = \tilde{\pi}(y)$$

so therefore $\tilde{\pi}P = \tilde{\pi}$.

We normalize this with $\sum_{x\in\mathcal{X}} \tilde{\pi}(x)$ to get a probability distribution. Notice that this normalization does not violate the stationarity condition. This is because scaling $\tilde{\pi}$ by any scalar c still satisfies the stationarity condition i.e if $\tilde{\pi}P = \tilde{\pi}$ and $\pi = c \cdot \tilde{\pi}$ then $\tilde{\pi}P = \tilde{\pi} \Rightarrow \pi P = \pi$. However, there is only one such c that makes π a probability distribution.

To complete the proof, it remains to show that $\sum_{x\in\mathcal{X}} \tilde{\pi}(x) = \mathbb{E}_z[\tau_z^+]$. This is shown below:

$$\begin{aligned}
\sum_{x\in\mathcal{X}} \tilde{\pi}(x) &= \sum_{x\in\mathcal{X}} \sum_{t=0}^{\infty} \Pr_z[X_t = x, \tau_z^+ > t] \\
&= \sum_{t=0}^{\infty} \sum_{x\in\mathcal{X}} \Pr_z[X_t = x, \tau_z^+ > t] && \text{(switching order of summations)} \\
&= \sum_{t=0}^{\infty} \sum_{x\in\mathcal{X}} \Pr_z[\tau_z^+ > t] \cdot \Pr_z[X_t = x] && \text{(independent events)} \\
&= \sum_{t=0}^{\infty} \Pr_z[\tau_z^+ > t] \\
&= \mathbb{E}_z[\tau_z^+]
\end{aligned}$$

Since $\sum_{x\in\mathcal{X}} \tilde{\pi}(x) = \mathbb{E}_z[\tau_z^+]$ (shown below), $\pi = \frac{\tilde{\pi}}{\mathbb{E}_z[\tau_z^+]}$ is a stationary distribution. ☐

***

We next argue that the stationary distribution is unique.

3 Proof of Uniqueness

Proposition 2. (Proof of uniqueness) The stationary distribution π of a finite, irreducible and aperiodic Markov Chain is unique.

Proof. Assume two distributions π_1 and π_2 that are both stationary distributions over $\mathcal{X}$. Consider a state $x \in \mathcal{X}$ that minimizes the ratio $\frac{\pi_1(x)}{\pi_2(x)}$. Let us call this minimizing value M. Then,

$$\pi_1(x) = \sum_{y\in\mathcal{X}} \pi_1(y) \cdot P(y,x) \geq \sum_{y\in\mathcal{X}} M \cdot \pi_2(y) \cdot P(y,x) = M \cdot \pi_2(y)$$

But we know that $\pi_1(x) = M \cdot \pi_2(x)$, so none of the $\pi_1(y) \geq M \cdot \pi_2(y)$ can be strict inequalities, i.e they must all be equalities. Therefore

$$\frac{\pi_1(x)}{\pi_2(x)} = M \text{ for all } x \in \mathcal{X}.$$

The following expansion shows that $M = 1$:

$$\begin{aligned}
\sum_{x\in\mathcal{X}} \pi_1(x) &= 1 && \text{(Sum of a probability distribution)} \\
\sum_{x\in\mathcal{X}} \frac{\pi_1(x)}{\pi_2(x)} \cdot \pi_2(x) &= 1 \\
\sum_{x\in\mathcal{X}} M \cdot \pi_2(x) &= 1 \\
M \sum_{x\in\mathcal{X}} \pi_2(x) &= 1 \\
M &= 1
\end{aligned}$$

Since $\pi_1(x) = \pi_2(x)$ for all $x \in \mathcal{X}$ (i.e $\pi_1 = \pi_2$), stationary distributions must be unique. ☐

4 π as the reciprocal of expected first return time

Finally, our work above has left us with this handy identity:

Proposition 3. For the stationary distribution π of a finite irreducible and aperiodic Markov Chain,

$$\pi(z) = \frac{1}{\mathbb{E}_z[\tau_z^+]}$$

Proof. We saw that

$$\pi(z) = \frac{\tilde{\pi}(z)}{\mathbb{E}_z[\tau_z^+]}$$

Since $\tilde{\pi}(z) = 1$ (because the number of visits a random walk of the chain makes to z before returning to z is exactly 1 - and this gets logged at the exact moment the random walk returns to z), the proposition follows immediately. □

5 Thinking through the assumptions

In this section we revisit the assumptions we made, and recall why they were needed.

Recall that we have proved the existence and uniqueness of a stationary distribution for a *finite, irreducible* and *aperiodic* Markov Chain.

- *finite*: What if our chain was not finite? We couldn't formulate a transition matrix P, which is necessary for our definition of irreducibility used in Lemma 6.1 makes use of this matrix P. I do not know if this assumption can be removed in a general setting.
- *irreducible*: We explicitly need irreducibility for Lemma 6.1 for proof of existence.

  More intuitively, consider this example for why a unique stationary distribution does not exist for reducible chains. Say A and B are disjoint subsets of $\mathcal{X}$ and $A \cup B = \mathcal{X}$. Then a sample path with $X_0 \in A$ would never get to B, and therefore could never "learn" the probability distribution of B. The stationary distribution is the long-term limiting distribution of the chain, and it should not depend where you start your chain from - but in the reducible case it is clear that the limiting distribution depends on where the chain starts. If we were to use our proofs above and were to even ignore the need for Lemma 6.1, we could show the existence of stationary distributions over a subgraph (those distributions would only have probability mass over either A or B, depending on where X_0 is). But we could not show uniqueness - indeed there would be two different stationary distributions, one each over A and B.
- *aperiodic*: Let us consider a periodic chain with two states A and B and $X_0 = A$ with period 2. Then, the probability distribution at all time steps when t is odd is $\{\Pr[X_t = B] = 1, \Pr[X_t = A] = 0\}$ and when t is even is $\{\Pr[X_t = B] = 0, \Pr[X_t = A] = 1\}$. This means that the probability distribution never converges - we do not have a limiting distribution. Therefore a stationary distribution does not exist for periodic chains in general.

References

[1] David A. Levin and Yuval Peres. *Markov Chains and Mixing Times.* 2nd ed. American Mathematical Society, 2017.

6 Appendix

Lemma 6.1. Let Y be a non-negative integer-valued random variable. Then,

$$\mathbb{E}[Y] = \sum_{t=0}^{\infty} \Pr[Y > t]$$

Proof. We first notice that

$$\mathbb{E}[Y] = \sum_{t=0}^{\infty} t \cdot \Pr[Y = t] = \sum_{t=1}^{\infty} t \cdot \Pr[Y = t]$$

We next notice that this final term can be expanded into a double sum:

$$\begin{aligned}
\sum_{t=1}^{\infty} t \cdot \Pr[Y = t] = \Pr[Y = 1] \\
&+ \Pr[Y = 2] + \Pr[Y = 2] \\
&+ \Pr[Y = 3] + \Pr[Y = 3] + \Pr[Y = 3] \\
&+ \Pr[Y = 4] + \Pr[Y = 4] + \Pr[Y = 4] + \Pr[Y = 4] \\
&+ \cdots \\
&\cdots \\
&= \left(\sum_{t=1}^{\infty} \Pr[Y = t]\right) + \left(\sum_{t=2}^{\infty} \Pr[Y = t]\right) + \left(\sum_{t=3}^{\infty} \Pr[Y = t]\right) + \cdots \\
&= \sum_{t=1}^{\infty} \sum_{i=t}^{\infty} \Pr[Y = i] \qquad (1)
\end{aligned}$$

Using the fact that $\Pr[Y > t] = \sum_{i=t+1}^{\infty} \Pr[Y = i]$, we can expand $\sum_{t=1}^{\infty} t \cdot \Pr[Y = t]$:

$$\begin{aligned}
\sum_{t=1}^{\infty} t \cdot \Pr[Y = t] &= \sum_{t=0}^{\infty} \sum_{i=t+1}^{\infty} \Pr[Y = i] \\
&= \sum_{t=1}^{\infty} \sum_{i=t}^{\infty} \Pr[Y = i] \qquad (2)
\end{aligned}$$

Combining (1) and (2) completes the proof. $\square$

Lemma 6.2. For any states $x, y \in \mathcal{X}$ of an irreducible Markov Chain,

$$\mathbb{E}_x[\tau_y^+] < \infty$$

Proof. We first use lemma (6.1) to observe that:

$$\mathbb{E}_x[\tau_y^+] = \sum_{t=0}^{\infty} \Pr_x[\tau_y^+ > t]$$

This motivates trying to bound $\Pr_x[\tau_y^+ > t]$.

By irreducibility, we know that there exists an integer $r > 0$ and a real number $\varepsilon > 0$ s.t for any $x, y \in \mathcal{X}$, there exists a $j \leq r$ s.t $P^j(x, y) > \varepsilon$. In other words. we know there exists an integer $j \leq r$ s.t the probability of a random walk starting at x and hitting y in j steps is $> \varepsilon$.

Then, the probability that the random walk does NOT hit y within r steps is $< (1 - \varepsilon)$, i.e

$$\Pr_x[\tau_y^+ > r] < (1 - \varepsilon)$$

Generalizing this argument, we see that for $k \geq 1$:

$$\Pr_x[\tau_y^+ > kr] < (1-\varepsilon)^k \tag{3}$$

We now expand our earlier $\mathbb{E}_x[\tau_y^+]$ to into sums of r terms

$$\begin{aligned}
\mathbb{E}_x[\tau_y^+] &= \sum_{t=0}^{\infty} \Pr_x[\tau_y^+ > t] \\
&= \sum_{t \in \{0,\cdots,r-1\}} \Pr_x[\tau_y^+ > t] \\
&+ \sum_{t \in \{r,\cdots,2r-1\}} \Pr_x[\tau_y^+ > t] \\
&+ \sum_{t \in \{2r,\cdots,3r-1\}} \Pr_x[\tau_y^+ > t] \\
&+ \cdots
\end{aligned}$$

Notice that the first sum in the above expansion can be bounded as

$$\sum_{t \in \{0,\cdots,r-1\}} \Pr_x[\tau_y^+ > t] \leq r \cdot \Pr_x[\tau_y^+ > 0]$$

This is because $\Pr_x[\tau_y^+ > t]$ is a decreasing function of t, which means that $\Pr_x[\tau_y^+ > 0]$ is the largest term in the above sum. Similarly,

$$\sum_{t \in \{r,\cdots,2r-1\}} \Pr_x[\tau_y^+ > t] \leq r \cdot \Pr_x[\tau_y^+ > r]$$

$$\sum_{t \in \{2r,\cdots,3r-1\}} \Pr_x[\tau_y^+ > t] \leq r \cdot \Pr_x[\tau_y^+ > 2r]$$

And in general,

$$\sum_{t \in \{kr,\cdots,kr-1\}} \Pr_x[\tau_y^+ > t] \leq r \cdot \Pr_x[\tau_y^+ > kr]$$

We can use this to get a general upper bound on $\mathbb{E}_x[\tau_y^+]$:

$$\begin{aligned}
\mathbb{E}_x[\tau_y^+] &\leq r \cdot \sum_{k=0}^{\infty} \Pr_x[\tau_y^+ > kr] \\
&\leq r \cdot \sum_{k=0}^{\infty} (1-\varepsilon)^k && \text{(using (3))}
\end{aligned}$$

Since $\sum_{k=0}^{\infty}(1-\varepsilon)^k < \infty$, the proof is complete. $\square$