Stationary distributions

We will now focus on Markov chains which are both irreducible and aperiodic and see how the distribution of X_n, here denoted $\mu^{(n)}$, changes as n grows larger.

We are interested in so called stationary distributions, that is, distributions expressing an equlibrium state for the chain. There will still be randomness, so X_n will change infinitely often, but the distribution of X_n will converge to some distribution, and this is the stationary distribution.

Though we are interested in the distribution and the convergence towards it, we must first establish the existence of such distributions. Given existence of a stationary distributions π, and that the distribution of X_n do converge to π we may wonder if there are more than one possibility for π. The uniqueness theorem states that π is unique for a chain, so if we can find one candidate, we have found the only one.

Theorems 5.1, 5.2 and 5.3 do the work for us, all that remain is to understand ...

- ...why stationary distributions exists.
- ...why convergence occur.
- ...why there only can be one.

The theorems give us the results, the proofs tells us why ! The proof of theorem 5.1 (existence) is useful when considering consequences of the uniqueness theorem. We start with two new concept ...

Hitting times for a Markov chain

Consider the Markov chain $(X_0, X_1, ...)$, and let $X_0 = s_i$, for some fixed but arbitrary state $s_i \in S$. Let

$$T_{i,j} = \min\{n \geq 1 : X_n = s_j\}$$

with $T_{i,j} = \infty$ if the chain never hits s_j. Let $\tau_{i,j} = \mathbb{E}[T_{i,j}]$ be the mean hitting time. The hitting time $T_{i,j}$ is the time it takes to reach state s_j when starting from s_i.

Lemma 5.1
For any irreducible aperiodic Markov chain with state space $S = \{s_1, ..., s_k\}$ and transition matrix P, we have for any two states $s_i, s_j \in S$ that if the chain starts in s_i then

$$\mathbb{P}(T_{i,j} < \infty) = 1.$$

Moreover, the mean hitting time $\tau_{i,j}$ is finite, i.e.

$$\mathbb{E}[T_{i,j} < \infty] = 1.$$

Convergence of distributions

When studying how distributions converge we need some sort of measure of how different distributions are.

Definition 5.1 ***Total variation***
If $\nu^{(1)} = (\nu_1^{(1)}, ..., \nu_k^{(1)})$ and $\nu^{(2)} = (\nu_1^{(2)}, ..., \nu_k^{(2)})$ are probability distributions on $S = \{s_1, ..., s_k\}$, then we define the total variation distance between $\nu^{(1)}$ and $\nu^{(2)}$ as follows.

$$d_{\text{TV}}(\nu^{(1)}, \nu^{(2)}) = \frac{1}{2}\sum_{l=1}^{k} |\nu_i^{(1)} - \nu_i^{(2)}|$$

If $\nu^{(1)}, \nu^{(2)}, \nu^{(3)}, ...$ and ν are probability distributions on S we say that $\nu^{(n)}$ converges to ν in total variation as $n \to \infty$, writing $\nu^{(n)} \overset{TV}{\to} \nu$, if

$$\lim_{n\to\infty} d_{\text{TV}}(\nu^{(n)}, \nu) = 0$$

Another, equivalent way to define total variation is the following.

$$d_{\text{TV}}(\nu^{(1)}, \nu^{(2)}) = \max_{A\subseteq S} |\nu^{(1)}(A) - \nu^{(2)}(A)|$$

Considering two different distributions, total variations give us the maximum difference between the probabilities they assign to any event. A consequence of the definition of the total variation distance is that

$$\nu^{(1)} = \nu^{(2)} \quad \Leftrightarrow \quad d_{\text{TV}}(\nu^{(1)}, \nu^{(2)}) = 0$$

for any two probability distribution $\nu^{(1)}$ and $\nu^{(2)}$ on S.

The theorems ...

Theorem 5.1 ***Existence of a stationary distribution***
For any irreducible and aperiodic Markov chain, there exists at least one stationary distribution.

Interpretation/Consequences:
For this class of Markov chains we have at last one stationary distribution, without it we cannot, in general, say much about the behaviour of $\mu^{(n)}$ as $n \to \infty$.

Idea of proof:
We make an educated guess, and see if the requirements for a stationary distribution are fulfilled andif they are we are done.

A resonable estimate of a stationary distribution is the following. Given $(X_0, ..., X_n)$ and any $s_i \in S$ let

$$\hat{\pi}(s_i) = \frac{1}{n} \sum_{k=0}^{n} I_{\{s_i\}}(X_n)$$

and as $n \to \infty$ the approximate probability $\hat{\pi}(s_i)$ converges to something as a consequence of the law of large numbers. When the chain returns to state s_i it starts all over again (due to the Markov property) so the realization of $(X_0, ..., X_n)$ as n grows larger, is just a repetition of $(X_0, ..., X_{T_{i,i}-1})$ where $T_{i,i}$ is the return time for state s_i.

Our educated guess for the stationary distribution is

$$\begin{aligned}
\pi(s_i) &= \frac{1}{\mathbb{E}[T_{1,1}]} \mathbb{E}\left[\sum_{k=0}^{T_{1,1}} I_{\{s_i\}}(X_n)\right] \\
&= \frac{1}{\mathbb{E}[T_{1,1}]} \sum_{k=0}^{T_{1,1}} \mathbb{E}[I_{\{s_i\}}(X_n)] \\
&= \frac{1}{\mathbb{E}[T_{1,1}]} \sum_{k=0}^{T_{1,1}} \mathbb{P}(X_n = s_i) \\
&= \frac{1}{\mathbb{E}[T_{1,1}]} \sum_{k=0}^{\infty} \mathbb{P}(X_n = s_i, n < T_{1,1})
\end{aligned}$$

where the numerator is the expected number of visits to state s_i between two visits to state s_1, and the denominator is the expected time we have between two such visits.

There is nothing special with state s_1 here, we could have choosen any state in S. For this to work we need both terms to be nonzero (trivial) and finite. Lemma 5.1 states that both hitting times and mean hitting times are finite for an irreducible and aperiodic chain.

What remain of the proof is to check that the requirements for a stationary distribution is fulfilled by our educated guess.

We state the theorem here since it plays a central role in the proof of theorem 5.3. The proof and its central ideas is postponed to next lecture.

Theorem 5.2 ***The Markov chain convergence theorem***
Let $(X_0, X_1, ...)$ be an irreducible aperiodic Markov chain with state space $S = \{s_1, ..., s_k\}$, transition matrix P, and arbitrary initial distribution $\mu^{(0)}$. Then, for any distribution π which is stationary for the transition matrix P, we have

$$\mu^{(n)} \stackrel{TV}{\to} \pi.$$

Equipped with convergence we are well prepared to deal with uniqueness.

Theorem 5.3 ***Uniqueness if the stationary distribution***
Any irreducible and aperiodic Markov chain has exactly one stationary distribution.

Interpretation/Consequences:
A consequence of this theorem and the proof of theorem 5.1 is that a stationary distribution for a irreducible aperiodic Markov chain is

$$\pi = \left(\frac{1}{\tau_{1,1}}, ..., \frac{1}{\tau_{k,k}}\right)$$

and this follows since in the proof of 5.1 we could have choosen any other state to treat as special. Then by uniqueness

$$\left(\frac{1}{\tau_{1,1}}, \frac{\rho_{1,2}}{\tau_{1,1}} ..., \frac{\rho_{1,k}}{\tau_{1,1}}\right) = \left(\frac{\rho_{2,1}}{\tau_{2,2}}, \frac{1}{\tau_{2,2}} ..., \frac{\rho_{2,k}}{\tau_{2,2}}\right) = ... = \left(\frac{\rho_{k,1}}{\tau_{k,2}}, \frac{\rho_{k,2}}{\tau_{k,k}} ..., \frac{1}{\tau_{k,k}}\right) = \left(\frac{1}{\tau_{1,1}}, ..., \frac{1}{\tau_{k,k}}\right)$$

since $\rho_{1,1} = \rho_{2,2} = ... = \rho_{k,k} = 1$ and

$$\frac{1}{\tau_{l,l}} = \frac{\rho_{k,l}}{\tau_{k,k}} \ , \quad \text{all } l, k \in \{1, ..., k\}$$

where $\rho_{j,i}$ is the number of visits to state s_j between two visits to s_i, note that $\rho_{1,k}$ and $\rho_{2,k}$ in general are different.

Idea of proof:
We use theorem 5.1 to conclude that there exists at least one stationary distribution, and the assume that there are at least two. We then apply theorem 5.2 and conclude that the two distributions are the same.

We let π and π' be two possibly different stationary distributions for a markov chain. Let the initial distribution $\mu^{(0)} = \pi'$, the since it is stationary $\mu^{(n)} = \pi'$ for every $n \geq 0$. Theorem 5.2 states that

$$\lim_{n\to\infty} d_{TV}(\mu^{(n)}, \pi) = \lim_{n\to\infty} d_{TV}(\pi', \pi) = 0$$

and since the last expression does not depend on n we have $d_{TV}(\pi', \pi) = 0$ and we are done.