

Continuous Time Markov Chains

EECS 126 (UC Berkeley)

Fall 2020

1 Introduction and Motivation

After spending some time with Markov Chains as we have, a natural question one might ask is what to do when our transitions between states don't come at nice, discrete, fixed intervals? For example, if we are modelling the number of people in line at checkout, it's not exactly reasonable to model this situation with customers arriving or being served with some probability exactly once every minute, or once every 10 minutes, or however long you want to make your interval. Instead, in the context of **Continuous Time Markov Chains**, we operate under the assumption that movements between states are quantified by *rates* corresponding to independent exponential distributions, rather than independent probabilities as was the case in the context of DTMCs.

2 Intuition and Building Useful Ideas

From discrete-time Markov chains, we understand the process of jumping from state to state. For each state in the chain, we know the probabilities of transitioning to each other state, so at each timestep, we pick a new state from that distribution, move to that, and repeat.

The new aspect of this in continuous time is that we don't necessarily know *when* a jump is going to happen: in addition to a distribution over the states, we need to provide some natural way to describe the distribution of the time you spend in each state. Ideally, we also want the machinery of CTMCs to be similar to what we've already developed for DTMCs.

One possibility might be to say each edge has to carry two pieces of information: a probability of making that transition, and a random variable describing how long that transition would take.




Figure 1: A possible (wrong) way for CTMCs to work

However, this model is too general and doesn't capture all the features of Markov chains that we want. It's also overly complicated; in addition to a transition matrix, we need to specify one random variable (not just a constant) per edge, so it would be difficult to have any of the discrete-time algebraic methods and definitions carry over. Therefore, let's look at the features we'd like to have, and see if a more natural description emerges.

2.1 Markov Property

The most essential property we'd like to carry over from DTMCs to CTMCs is the **Markov property**. Informally, this says that wherever we end up next is only dependent on where we are right now, rather than the entire history of where we have been. Mathematically, the property is:

$$\mathbb{P}(X_{t+1} = x_{t+1} | X_t = x_t, X_{t-1}, \dots, X_0 = x_0) = \mathbb{P}(X_{t+1} = x_{t+1} | X_t = x_t).$$

In continuous time, we can extend this slightly: not only do we want independence from the whole history of where we've been, we want independence from *how long* we've been where we are right now. If a CTMC is at state 1, it doesn't matter if the chain just moved there or if it's been there for some time (five minutes, five hours, etc.); the time required to leave the state should only depend on the chain setup and not on the history of what's happened.

2.2 Why Jump Times are Exponential

Although the transition times are no longer fixed to be 1, we can use the Markov property to figure out how they are distributed. Suppose at time t , we're in state i , and we're interested in the distribution of τ , the time until the chain jumps to a different state. As we said above, a key property of τ is that it's independent of how much time we have *already* spent at i . That is, if the chain has already spent time T in its current state, the distribution of τ is exactly the unconditioned distribution of $T + \tau$:

$$f_{\tau|\tau>T}(t + T) = f_{\tau}(t)$$

This is the statement that τ is memoryless, and we know that the unique continuous memoryless distribution is the exponential! Therefore, $\tau \sim \text{Exponential}(q)$. We will see that not only are the transition times at each state exponentially distributed, but transitions between specific pairs of states are also described by a rate of an exponential random variable. The value of q is dependent on each pair of states we're interested in (we conventionally denote the rates by q_{ij} for a transition $i \rightarrow j$), and is encoded in the rate matrix Q .

2.3 Intuition per Step

We just saw that the jump times are exponential, and so it makes sense to describe how fast the transition $i \rightarrow j$ is by the parameter q_{ij} of the exponential random variable that represents its jump time. When we draw a CTMC, these parameters are the edge weights, so now we can make a more useful version of Figure 1:




Figure 2: A corrected CTMC diagram

From this diagram, how do we get the probability of transitions? Intuitively, we now know how fast each transition happens: for example, in Figure 2, the parameter associated with $1 \rightarrow 2$ is 6 and that associated with $1 \rightarrow 3$ is 2. Therefore, transitions $1 \rightarrow 2$ happen three times as fast as those $1 \rightarrow 3$. This means that if we start at 1, then $3/4$ of the time we should expect our next state to be 2, and the remaining $1/4$ of the time it's 3. We can formalize this!

Recall that the minimum of exponential RVs is also exponential, with parameter equal to the sum of the rates. This gives us two high-level ways of looking at CTMCs. Suppose from our current state we have n possible states, $1, 2, \dots, n$, and their exponential rates are $\lambda_1, \lambda_2, \dots, \lambda_n$ (that is, the time required to jump to state 1 is $\text{Exponential}(\lambda_1)$, the time required to jump to state 2 is $\text{Exponential}(\lambda_2)$, and so on). Then, we can view the action of a CTMC as either of the following:

1. Draw proposed jump times $\tau_1 \sim \text{Exponential}(\lambda_1)$, $\tau_2 \sim \text{Exponential}(\lambda_2)$, \dots , $\tau_n \sim \text{Exponential}(\lambda_n)$ and jump to the state that comes up first.
2. Draw a jump time $\tau \sim \text{Exponential}(\lambda_1 + \lambda_2 + \dots + \lambda_n)$, wait that much time, and jump to a state from the distribution given by $\mathbb{P}(X_j = k) = \frac{\lambda_k}{\sum_i \lambda_i}$.

This also tells us that the time that we stay put is distributed according to $\text{Exponential}(\sum_i \lambda_i)$, which means the mean time is $\frac{1}{\lambda_1 + \lambda_2 + \dots + \lambda_n}$: the higher all the outgoing rates are, the less time before we jump.

2.4 Self-Loops

You may have noticed that Figure 2 lacks self-loops, i.e. the idea from DTMCs of a state transitioning to itself. However, since we don't have fixed time steps any more, this concept does not translate to CTMCs. Since we're not 'forced' to ever move, a state transitioning to itself is just the same as nothing happening.

However, the timescale on which 'nothing' happens is still meaningful: as we just saw, the time for which we stay in the state we're currently in is distributed according to $\text{Exponential}(\sum_i \lambda_i)$, and so the rate at which we transition from a state to any other state is $\sum_i \lambda_i$. In the rate matrix, we'll denote this by $-\sum_i \lambda_i$, to indicate that it's the rate of *entering* a state rather than *leaving* it.

3 Rate Matrix, Stationary Distributions, and Hitting Times

3.1 Basic Definitions

We now define what exactly a CTMC is. Concretely, we start with some set of states \mathcal{X} , along with some initial probability distribution over \mathcal{X} . Our random process is then defined by a **rate matrix** $Q \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{X}|}$ where

$$\begin{aligned} Q(i, j) &\geq 0 \quad \forall i \neq j \\ \sum_j Q(i, j) &= 0 \end{aligned}$$

The above in particular implies that $Q(i, i) = -\sum_{j \neq i} Q(i, j)$.

Example 1. We could have

$$Q = \begin{bmatrix} -4 & 3 & 1 \\ 0 & -2 & 2 \\ 1 & 1 & -2 \end{bmatrix},$$

and this would be a perfectly valid rate matrix for a CTMC with $|\mathcal{X}| = 3$ (three states). (Note: This follows the convention in Walrand's book, not Bertsekas).

The associated CTMC is shown in Figure 3.




Figure 3: The CTMC with rate matrix Q

More generally, our rate matrix Q looks like:

$$Q = \begin{bmatrix} -\sum_{i=1, i \neq 1}^n \lambda_{1 \rightarrow i} & \lambda_{1 \rightarrow 2} & \dots & \lambda_{1 \rightarrow n} \\ \lambda_{2 \rightarrow 1} & -\sum_{i=1, i \neq 2}^n \lambda_{2 \rightarrow i} & \dots & \lambda_{2 \rightarrow n} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{n \rightarrow 1} & \lambda_{n \rightarrow 2} & \dots & -\sum_{i=1, i \neq n}^n \lambda_{n \rightarrow i} \end{bmatrix},$$

where each λ represents the rate at which we follow a particular edge, which is an exponential random variable (i.e. the first arrival of a Poisson process, if you're familiar with that). For simplicity, we denote

$$q_i = \sum_{j=1, j \neq i}^n \lambda_{i \rightarrow j} = -Q(i, i)$$

Where the last equality follows from the definition of the rate matrix. Let's make sure our setup does actually satisfy the Markov property, at least at a small enough time scale. To see why the continuous time analog we have constructed might satisfy something like this property we have to first recall an important property of exponentials. Suppose $Y \sim \text{Exponential}(\lambda)$. Then we have that

$$\mathbb{P}(Y \geq \epsilon) = e^{-\lambda\epsilon} = 1 - \lambda\epsilon + o(\epsilon),$$

where we have used the Taylor series expansion of e^x and used the "little-o" notation, which just refers the set of functions such that:

$$\lim_{\epsilon \rightarrow 0} \frac{o(\epsilon)}{\epsilon} = 0.$$

For example, $\epsilon^2 \in o(\epsilon)$. Why does this property of exponential RVs matter? Suppose we are at state i with holding rate q_i . Then we have

1. no transitions happen with probability $1 - q_i\epsilon + o(\epsilon)$ or
2. 1 transition happens with probability $q_i\epsilon + o(\epsilon)$
3. 2 or more transitions happen with *negligible* probability $o(\epsilon)$.

In the context of our CTMCs, it means we have the following property:

Proposition 1.

$$\mathbb{P}(X_{t+\epsilon} = j | X_t = i; X_u, u < t) = \begin{cases} \epsilon Q(i, j) + o(\epsilon) & i \neq j \\ 1 + \epsilon Q(i, i) + o(\epsilon) & i = j \end{cases}$$

What the above tells us is that our CTMCs satisfy the Markov property, provided ϵ is small enough. Why is the above true? We saw above that for small enough ϵ , the probability that we observe two or more "jumps" or "transitions" in our process is negligibly small. Also recall that if we have independent $Y_1 \sim \text{Exponential}(\lambda_1), \dots, Y_n \sim \text{Exponential}(\lambda_n)$, then

$$\min_i Y_i \sim \text{Exponential}\left(\sum_{i=1}^n \lambda_i\right).$$

This tells us that our **holding time** at state i , or the amount of time we wait before transitioning to a new state from state i , is also exponentially distributed with rate $\lambda = \sum_{j \neq i} Q(i, j) = -Q(i, i)$. Recall that we also proved in the homework that

$$\mathbb{P}(\min_i Y_i = Y_j) = \frac{\lambda_j}{\sum_i \lambda_i}.$$

The above facts tell us that

$$\begin{aligned} \mathbb{P}(X_{t+\epsilon} = j | X_t = i; X_u, u < t) &= \mathbb{P}(\text{we leave state } i) \mathbb{P}(\text{we go to state } j | \text{we leave state } i) \\ &\approx \epsilon(-Q(i, i)) \cdot \frac{Q(i, j)}{-Q(i, i)} = \epsilon Q(i, j), \end{aligned}$$

where in the above we have omitted the $o(\epsilon)$ terms for clarity. We conclude this section by noting that a Poisson process can be modeled very simply as a CTMC:

Remark 1. If we have a $PP(\lambda)$, we can model it as a CTMC as in Figure 4.




Figure 4: Poisson process and CTMC connection

3.2 Stationary Distributions and Hitting Times

Now we are ready to discuss stationary distributions. Recall that for DTMCs, we were interested in solving for π such that $\pi = \pi P$, where P was our transition probability matrix. However, in the context of CTMCs, we instead need to solve for

$$\pi Q = 0, \quad \sum_i \pi_i = 1.$$

One might wonder, why the change? For DTMCs, the whole reason we formulated the original equation was just to encode the notion that the probability mass going into and coming out of any given state i remains consistent when we are in the stationary distribution:

$$\underbrace{\pi_i}_{\text{prob mass out of } i} = \underbrace{\sum_{j \in \mathcal{X}} P(j, i) \pi_j}_{\text{prob mass into } i}$$


Here, we don't have probabilities, but rather want to make sure that the combined rates going into a state equals the combined rate going out of the state. Intuitively, this will ensure that when our chain is in stationarity, the amount of probability mass that remains in any given state remains roughly the same:

$$\forall i, \quad \underbrace{\pi_i \cdot (-Q(i, i))}_{\text{rate out}} = \underbrace{\sum_{j \neq i} \lambda_{j \rightarrow i} \pi_j}_{\text{rate in}}$$

$$\iff \pi Q = 0.$$

Therefore, solving for our stationary distribution just boils down to solving $\pi Q = 0$.

Exercise 1 (B&T 7.14). Say our application has a normal state (1) and a test state (2) and a repair state (3).



What is the stationary distribution of this Markov Chain? Note we have

$$Q = \begin{bmatrix} -1 & 1 & 0 \\ 2.5 & -5 & 2.5 \\ 3 & 0 & -3 \end{bmatrix}$$

Write out the balance equations, either by examining $\pi Q = 0$ or by just looking at the "rate in = rate out" for each state. Show that

$$\pi = (30/41, 6/41, 5/41).$$

The last fundamental thing we have to worry about is solving hitting time problems using first step equations. It is very analagous to solving hitting time problems with DTMCs. Can you guess what the differences will be? Recall that when we set up our FSEs in the context of DTMCs, we had to remember to add 1 to every equation to take into account our discrete time intervals going by. In the context of CTMCs, this no longer makes sense, and instead we have to consider the *expected* holding time we stay in a given state before jumping. This is exactly $\frac{1}{-Q(i,i)} = \frac{1}{q_i}$. This concept is best illustrated by example:

Example 2. Consider 20 lightbulbs that have indep. lifetimes that are exponentially distributed with rate 1 (month). How long before all the bulbs die out?

We can model this with 21 states, corresponding to the number of lightbulbs still alive. We transition from state 20 to state 19 with rate 20 (min of 20

exponentials!), and from 19 to 18 with rate 19, and so on (draw the Markov chain if it's not clear!). How long does it take for us to get to zero? We have to come up with first step equations. The biggest difference is that we are not transitioning after 1 time step now, we are transitioning on average after $1/q_i$ amount of time (the mean of the exponential) if we are in state i . We can define x_i as the expected amount of time to hit zero given you are in state i . We have then

$$x_i = \frac{1}{i} + x_{i-1}$$

This implies that

$$x_{20} = \frac{1}{20} + \frac{1}{19} + \dots + 1 \approx 3.6$$

Example 3. Now assume the burnt out bulbs are replaced after an exponential amount of time with mean 0.1 month (meaning $\lambda = 10$). Now what is the expected amount of time until all our bulbs have burnt out? Now our FSE looks like:

$$\beta(20) = \frac{1}{20} + \beta(19)$$

and more generally for $1 \leq m \leq 19$:

$$\beta(m) = \frac{1}{m+10} + \frac{m}{m+10}\beta(m-1) + \frac{10}{m+10}\beta(m+1)$$

and finally $\beta(0) = 0$. Solving these equations recursively yields

$$\beta(20) \approx 2488$$

4 Jump Chains and Connection to DTMCs

As we have already mentioned, there is a close connection between CTMCs and DTMCs. In this section, we explore this connection a bit more concretely and clear up some common misconceptions. There are two ways to "model" a CTMC as a DTMC, each of which comes with their own assumptions and exhibits different behaviors. We will start with what is known in the business as a **jump** or **embedded** (or even **embedded jump**) chain. Intuitively, a jump chain captures exactly what our sequence of random variables (X_1, \dots) would look like if we ignored the holding time altogether, and simply considered the next state we jump to. Mathematically, we define a jump chain to have

the transition probabilities $P(i, j) = \frac{Q(i, j)}{q_i}$ and $P(i, i) = 0$. Notice, by the last equation, there are no self loops, while the first equation captures exactly the probability that, given we make a jump out of state i , that jump actually takes us to state j . A natural question to ask ourselves is whether or not this DTMC has the same stationary distribution as its corresponding CTMC. The answer is no, and is illustrated by the following example.

Example 4. We can consider the two state markov chain which transitions from state 0 to state one with rate λ , and from state 1 to state 0 with rate μ . Then by the balance equations we have

$$\begin{aligned}\lambda\pi_0 &= \mu\pi_1 \\ \pi_0 + \pi_1 &= 1\end{aligned}$$

We can solve these equations to find that

$$\pi_0 = \frac{\mu}{\lambda + \mu} \quad \pi_1 = \frac{\lambda}{\lambda + \mu}$$

So when $\lambda = 1$ and $\mu = 2$, we are "parked" in state 0 twice as much as we are parked in state 1. However, the embedded DTMC has the stationary distribution $(1/2, 1/2)$, no matter what λ and μ are (**note**: the jump chain is periodic so it won't converge to this unless the initial distribution is also $(1/2, 1/2)$, but nonetheless $(1/2, 1/2)$ still satisfies the balance equations so it is a stationary distribution). Clearly, we need to do something different if we want our DTMC to have the same stationary distribution as the CTMC.

So somehow this "jump" chain is losing some information regarding the holding times about our CTMC. The next natural question to ask ourselves is whether we can relate the stationary distribution of a CTMC with its corresponding jump chain at all. It turns out the answer to this question is yes! Suppose π is a stationary distribution for our CTMC, and ψ the stationary distribution for our corresponding jump chain. Then we know that ψ satisfies

$$\forall i \in \mathcal{X}, \quad \psi_i = \sum_j \psi_j P(j, i),$$

and we know that π satisfies

$$\pi Q = 0 \iff \forall i \in \mathcal{X}, \quad \pi_i q_i = \sum_{j \neq i} \pi_j Q(j, i).$$

But, by definition of the jump chain, we also know that $Q(j, i) = P(j, i)q_j$. Note that if $i = j$, then $P(i, j) = 0$. Then plugging this into our stationary equation for π tells us that

$$\forall i \in \mathcal{X}, \quad \pi_i q_i = \sum_{j \neq i} \pi_j P(j, i) q_j.$$

Examining the above a bit more closely actually tells us that the vector $(\pi_i q_i)_{i \in \mathcal{X}}$ is an eigenvector for P with eigenvalue 1, meaning if we normalize it then it is a valid stationary distribution for our jump chain! That is,

$$\psi_i = \frac{\pi_i q_i}{\sum_j \pi_j q_j}$$

is our stationary distribution for our jump chain, in terms of the stationary distribution of our CTMC. Likewise, given the stationary distribution of a jump chain ψ , one can compute the stationary distribution for the CTMC as

$$\pi_i = \frac{\psi_i / q_i}{\sum_j \psi_j / q_j}.$$

That's all very nice, but what if we want to accurately simulate our CTMC with a DTMC? That is, is it possible to construct a DTMC with the same stationary distribution as its corresponding CTMC? The answer is again yes, and our guiding intuition will be to somehow incorporate the holding times into our DTMC that the jump chain so woefully neglected. The way to do this is simple: add self loops where they are necessary! Simply define some q , and let $P(i, j) = \frac{Q(i, j)}{q}$, and then our self loop, instead of being zero, will be $P(i, i) = 1 - \frac{q_i}{q}$. We must have that our probabilities are all nonnegative, so we pick $q \geq \max_i q_i$. Written in matrix form, we are letting

$$P = I + \frac{1}{q}Q.$$

We can verify that any stationary distribution for P is also a stationary distribution for Q , and vice versa:

$$\begin{aligned} \pi P = \pi &\iff \pi \left(I + \frac{1}{q}Q \right) = \pi \\ &\iff \pi + \pi \frac{1}{q}Q = \pi \end{aligned}$$

$$\iff \pi Q = 0.$$

This technique is sometimes referred to as "uniformization", but the term itself is not so important. It is more important that you remember what the fundamental difference between this DTMC and the jump DTMC is.

5 Kolmogorov Equations

In discrete time, the Chapman-Kolmogorov equations encoded a useful fact about how Markov chains worked over multiple timesteps:

$$\mathbb{P}(X_{n+m} = j | X_0 = i) = \sum_{k \in \mathcal{X}} \mathbb{P}(X_{n+m} = j | X_n = k) \mathbb{P}(X_n = k | X_0 = i)$$

In words, this says the probability of reaching a certain result in n steps is equal to the sum of the probabilities of each path to get there.

A similar equation is true in continuous time, but here it'll allow us to write a differential equation that lets us work with a CTMC rate matrix more easily. In continuous time, let s and t both be time intervals; then,

$$\mathbb{P}(X_{s+t} = j | X_0 = i) = \sum_{k \in \mathcal{X}} \mathbb{P}(X_{s+t} = j | X_s = k) \mathbb{P}(X_s = k | X_0 = i).$$

This looks almost the same as the discrete-time result, but in the continuous-time setting we can do more with this. Consider the matrix $P(t)$ to be the transition probabilities in a fixed time interval t (following the same conventions as DTMCs: $P_{ij}(t)$ is the probability of transitioning $i \rightarrow j$ in time t). Then, the above can be more easily written as

$$P_{ij}(s+t) = \sum_{k \in \mathcal{X}} P_{ik}(s) P_{kj}(t),$$

In words, to get from $i \rightarrow j$ in time $s+t$ we have to move $i \rightarrow$ some k in time s and then move $k \rightarrow j$ in time t , over all the possible k . The above equation is the definition of matrix multiplication, so we can say

$$P(s+t) = P(s)P(t)$$

This is the analogue of the k -step discrete time transition matrix being the k th power of the transition matrix. We can relate this to the rate matrix by differentiating in time. Let $s \rightarrow \Delta t$ and apply the definition of a derivative:

$$\begin{aligned}
 P(t + \Delta t) &= P(t)P(\Delta t) \\
 P(t + \Delta t) - P(t) &= P(t)(P(\Delta t) - I) \\
 \frac{P(t + \Delta t) - P(t)}{\Delta t} &= P(t)\frac{P(\Delta t) - I}{\Delta t} \\
 P'(t) = \lim_{\Delta t \rightarrow 0} \frac{P(t + \Delta t) - P(t)}{\Delta t} &= P(t) \lim_{\Delta t \rightarrow 0} \frac{P(\Delta t) - I}{\Delta t}
 \end{aligned}$$

Intuitively, this derivative expression should be related to the rate matrix (both the derivative of the transition matrix and the rate matrix are supposed to encode how fast transitions happen), and we can show that in the limit of $\Delta t \rightarrow 0$, $Q = \frac{P(\Delta t) - I}{\Delta t}$.

First, we look at the off-diagonal terms. Recall from our rate matrix setup that for short times ϵ , the jump probabilities are $P(\Delta t)_{ij} = q_{ij}\Delta t$, so the off-diagonal terms of the limit $\lim_{\Delta t \rightarrow 0} \frac{P(\Delta t) - I}{\Delta t}$ must be q_{ij} . Next, we look at the diagonal terms. Since $P(\Delta t)$ is a Markov probability matrix, the diagonal terms must be set so that each row sums to 1, i.e.

$$P(\Delta t)_{ii} = 1 - \sum_{j=1}^n q_{ij}\Delta t$$

We're concerned with $\frac{P(\Delta t) - I}{\Delta t}$, so subtracting 1 and dividing by Δt gets us

$$\left(\frac{P(\Delta t) - I}{\Delta t}\right)_{ii} = - \sum_{j=1, j \neq i}^n q_{ij}.$$

This is precisely how we previously defined the diagonal elements of Q ! Therefore we can conclude that $\lim_{\Delta t \rightarrow 0} \frac{P(\Delta t) - I}{\Delta t} = Q$ and we have the differential equation

$$\frac{dP(t)}{dt} = P(t)Q \tag{1}$$

This is the *Kolmogorov forward equation*. Note that in some sources, Q is called the Markov chain's generator or infinitesimal generator and is denoted by A .

If we did this whole process instead starting with the equation $P(t + \Delta t) = P(\Delta t)P(t)$, we would get the differential equation

$$\frac{dP(t)}{dt} = QP(t) \tag{2}$$

This is the *Kolmogorov backward equation*.

Note that if P were a real number, these equations would be satisfied by an exponential function, so analogous to that, we define the *matrix exponential* by substituting tQ into the parameter for the exponential Taylor series:

$$e^{tQ} = \sum_{i=0}^{\infty} \frac{t^i Q^i}{i!}.$$

Therefore, the Kolmogorov equations give us a quick way of finding the transition matrix within a time t from the rate matrix:

$$P(t) = P(0)e^{tQ} = e^{tQ} \tag{3}$$

Solving this analytically requires taking the Jordan normal form of a matrix (beyond our scope), but if you're ever working with CTMCs in a computational setting there's usually inbuilt functions that can take matrix exponentials! See <http://www.cs.cornell.edu/cv/ResearchPDF/19ways+.pdf> for more.