Continuous Time Markov Chains (CTMC)

Start with real numbers (15 minutes)

Motivation

Imagine a toddler who has only two moods (states):

- Happy (State 0)
- Crying (State 1)

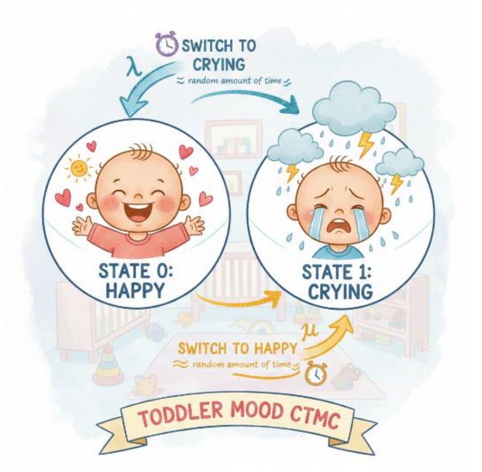

This is a CTMC because the toddler stays in one mood for a random amount of time before switching to the other.

The States

- State 0: Happy
- State 1: Crying

The Transition Rates

In continuous time, we don't use fixed probabilities (like "50% chance to switch"). Instead, we use **rates** (how fast something tends to happen).

Rate $\boldsymbol{\lambda}$: *The rate at which the toddler goes from **Happy → Crying**.*

Let's say $\boldsymbol{\lambda} = \mathbf{2}$ times per hour. (This means, on average, they get upset twice an hour).

Rate $\boldsymbol{\mu}$: *The rate at which the toddler goes from **Crying → Happy**.*

Let's say $\boldsymbol{\mu} = \mathbf{4}$ times per hour. (They recover quickly, usually within 15 minutes).

How it works in Continuous Time

- ***The Waiting Game:*** *If the toddler is currently **Happy**, they will stay happy for a random duration. In a CTMC, this duration follows an **Exponential Distribution**.*
- ***Memoryless:*** *It doesn't matter how long the toddler has already been happy. The probability of them crying in the next minute is always the same. They don't "get tired" of being happy; the switch happens randomly based on the rate λ .*

A Slightly More Complex Example: The Coffee Shop Queue (M/M/1 Queue)

This is the most famous example, often called the Birth-Death Process.

- States: The number of people in line $(0, 1, 2, 3, \dots)$.
- Transitions:
  - Arrival (λ): A new customer arrives, and the state goes up (0 → 1, or 1 → 2). This happens randomly (e.g., 10 customers per hour).
  - Service (μ): The barista finishes serving a customer, and the state goes down (2 → 1). This also happens randomly (e.g., the barista serves 12 customers per hour).

Why is this "Continuous"?

In a discrete model (like a turn-based game), a customer would arrive exactly every minute. In this Continuous model, a customer might arrive at 10:01, the next at 10:02:30, and the next at 10:15. The system evolves smoothly in real-time.

Continuous Time Markov Chain (CTMC)

A **Continuous-Time Markov Chain (CTMC)** is a stochastic process that evolves over a continuous time set (usually $[0, \infty)$) and satisfies the Markov property.

Formal Definition

A continuous-time Markov chain (CTMC) is a stochastic process

$$\{X(t), t \geq 0\}$$

with a countable state space S, satisfying the Markov property:

$$P(\, X(t+h) = j \mid X(t) = i, \text{ past history}\,) = P(\, X(t+h) = j \mid X(t) = i\,), \forall\, i, j \in S.$$

and whose transition behavior is characterized by a generator matrix $Q = [q_{ij}]$.

Steps to Construct the Generator Matrix Q

1. List the states:
Identify all possible states of the system and label them (e.g., $0,1,2, \ldots$).

2. Determine transition rates:
For each pair of *different* states $i \neq j$, find the rate q_{ij}.

3. Fill off-diagonal entries:
Place the transition rates in the matrix:
$q_{ij} =$ rate of transition from i to $j, i \neq j$.

4. Set impossible transitions to zero:
If a transition from i to j is not allowed, then set $q_{ij} = 0$.

5. Compute diagonal entries
For each state i, set the diagonal entry as:

$$q_{ii} = -\sum_{j \neq i} q_{ij}$$

(negative of the total rate of leaving state i).

6. Check row sums
Verify that each row of the Q-matrix sums to zero:

$$\sum_j q_{ij} = 0.$$

Finding the **stationary distribution (or stationary probabilities)** is *one important way* to study the long-term behavior of a Continuous-Time Markov Chain (CTMC)

Let's see the example of the Toddler, which was discussed above:

Example 1:

A toddler has two moods modelled by a continuous-time Markov chain (CTMC):

State 0: Happy, State 1: Crying

The toddler switches from Happy to Crying at rate 2 per hour, and from Crying to Happy at rate 4 per hour.

(a) Write the Q-matrix of the CTMC.

(b) Find the stationary distribution.

Step 1: Define the States and Rates

- **State 0:** Happy
- **State 1:** Crying

Step 2: Determine the transition rates

- **Rate 0 → 1:** $q_{01} = 2$ per hour (gets upset)
- **Rate 1 → 0:** $q_{10} = 4$ per hour (recovers)

Step 3: Construct the Generator Matrix Q

From state 0 (Happy):

- Can only go to state 1 at rate 1, that is

$$q_{01} = 2 \text{ and } q_{00} = -2 \text{ (negative total exit rate)}$$

From state 1 (Crying):

- Can only go to state 0 at rate 4, that is

$$q_{10} = 4 \text{ and } q_{11} = -4 \text{ (negative total exit rate)}$$

Generator Matrix:

$$Q = \begin{bmatrix} -2 & 2 \\ 4 & -4 \end{bmatrix}$$

Check: Row sums = 0.

Step 4: Set Up Stationary Equations

Let $\pi = [\pi_0 \quad \pi_1]$, that is

- π_0 = long-run probability of being in state 0 (Happy)
- π_1 = long-run probability of being in state 1 (Crying)

We solve $\pi Q = 0$ with $\pi_0 + \pi_1 = 1$.

From $\pi Q = 0$:

$$\begin{aligned} \pi_0 \cdot (-2) + \pi_1 \cdot 4 &= 0 \\ \pi_0 \cdot 2 + \pi_1 \cdot (-4) &= 0 \end{aligned}$$

These are the same equation. Using the first:

$$-2\pi_0 + 4\pi_1 = 0$$
$$2\pi_0 = 4\pi_1$$
$$\pi_0 = 2\pi_1$$

Step 4: Add Normalization Condition

$$\pi_0 + \pi_1 = 1$$

Substitute $\pi_0 = 2\pi_1$:

$$2\pi_1 + \pi_1 = 1$$
$$3\pi_1 = 1$$
$$\pi_1 = \frac{1}{3} \approx 0.333$$
$$\pi_0 = 2 \cdot \frac{1}{3} = \frac{2}{3} \approx 0.667$$

Step 5: Stationary Distribution

$$\pi = [\pi_0 \quad \pi_1] = \begin{bmatrix}\frac{2}{3} & \frac{1}{3}\end{bmatrix} \approx [0.667 \quad 0.333]$$

Interpretation:

- **66.7%** of time: Toddler is happy (state 0)
- **33.3%** of time: Toddler is crying (state 1)

Example 2:

A machine in a factory can be in one of three conditions:

- **State 0:** The machine is working normally
- **State 1:** The machine has a minor failure
- **State 2:** The machine has a major failure

The machine's condition changes over time according to the following rules:

- When the machine is working normally (State 0), it develops a **minor failure** at a rate of **2 per hour**.
- When the machine has a minor failure (State 1), it worsens into a **major failure** at a rate of **1 per hour**.
- When the machine has a major failure (State 2), it is repaired and returns to the **normal working state** at a rate of **3 per hour**.

Assume the machine behaviour can be modelled as a **continuous-time Markov chain**.

States:

A machine has 3 states:

- 0 = Working normally 1 = Minor failure 2 = Major failure

Rates:

- 0 → 1 at rate 2, i.e. $q_{01} = 2, q_{02} = 0.$
- 1 → 2 at rate 1, i.e. $q_{12} = 1, q_{10} = 0.$
- 2 → 0 after repair at rate 3, i.e. $q_{20} = 3, q_{21} = 0.$

Q-matrix:

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

This is a simple CTMC with 3 states and realistic interpretation.

2) Balance equations $\boldsymbol{\pi Q = 0}$

Take $\pi = [\pi_0 \quad \pi_1 \quad \pi_2]$, then

$$\pi Q = 0 \Rightarrow [\pi_0 \quad \pi_1 \quad \pi_2]\begin{bmatrix} -2 & 2 & 0 \\ 0 & -1 & 1 \\ 3 & 0 & -3 \end{bmatrix} = 0$$

$$\Rightarrow \begin{bmatrix} -2\pi_0 + 3\pi_2 \\ 2\pi_0 - \pi_1 \\ \pi_1 - 3\pi_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

Thus

$$-2\pi_0 + 3\pi_2 = 0 \Rightarrow 3\pi_2 = 2\pi_0 \Rightarrow \pi_2 = \frac{2}{3}\pi_0.$$

and

$$2\pi_0 - \pi_1 = 0 \Rightarrow \pi_1 = 2\pi_0.$$

3) Normalize $\sum \pi_i = 1$

$$\pi_0 + \pi_1 + \pi_2 = \pi_0 + 2\pi_0 + \frac{2}{3}\pi_0 = \frac{11}{3}\pi_0 = 1$$

so $\pi_0 = 3/11$.
Therefore

$$\boxed{\pi = (\pi_0, \pi_1, \pi_2) = \left(\frac{3}{11}, \frac{6}{11}, \frac{2}{11}\right)}$$

Interpretation

- Long-run fraction of time machine is **working**: $\pi_0 = 3/11 \approx 0.2727$.
- In **minor failure**: $6/11 \approx 0.5455$.
- In **major failure**: $2/11 \approx 0.1818$.

Example 3

During a study session in the university library, a student's activity changes randomly over time and can be modelled as a continuous-time Markov chain (CTMC) with three states:

Studying, Distracted and Taking a Break

The transition rates (per five minutes) are given as follows:

Studying → Distracted at rate 3	Studying → Taking a Break at rate 2
Distracted → Studying at rate 2	Distracted → Taking a Break at rate 1
Taking a Break → Studying at rate 1	Taking a Break → Distracted at rate 2

(a) Write the generator matrix (Q-matrix) of the CTMC.

(b) Find the stationary probabilities of the CTMC.

Solution:

State Space: Let the states be ordered as

0 = Studying,   1 = Distracted,   2 = Taking a Break.

Rates:

- 0 → 1 at rate 3, i.e. $q_{01} = 3$ and 0 → 2 at rate 2, i.e. $q_{02} = 2$.
- 1 → 0 at rate 2, i.e. $q_{10} = 2$ and 1 → 2 at rate 1, i.e. $q_{12} = 1$.
- 2 → 0 at rate 1, i.e. $q_{20} = 1$ and 2 → 1 at rate 2, i.e. $q_{21} = 2$.
- $q_{00} = -(q_{01} + q_{02}) = -5$. Also $q_{11} = -3$, and $q_{22} = -3$.

Generator matrix (Q-matrix):

Hence the generator matrix is

$$Q = \begin{bmatrix} -5 & 3 & 2 \\ 2 & -3 & 1 \\ 1 & 2 & -3 \end{bmatrix}.$$

Balance equations $\pi Q = 0$

Take $\pi = [\pi_0 \quad \pi_1 \quad \pi_2]$, then

$$\pi Q = 0 \Rightarrow [\pi_0 \quad \pi_1 \quad \pi_2]\begin{bmatrix} -5 & 3 & 2 \\ 2 & -3 & 1 \\ 1 & 2 & -3 \end{bmatrix} = 0$$

$$\Rightarrow \begin{bmatrix} -5\pi_0 + 2\pi_1 + \pi_2 \\ 3\pi_0 - 3\pi_1 + 2\pi_2 \\ 2\pi_1 + \pi_1 - 3\pi_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

This gives

$$-5\pi_0 + 2\pi_1 + \pi_2 = 0 - - - (i)$$
$$3\pi_0 - 3\pi_1 + 2\pi_2 = 0 - - - (ii)$$
$$2\pi_1 + \pi_1 - 3\pi_2 = 0 - - - (iii)$$

From (i):

$$\pi_2 = 5\pi_0 - 2\pi_1 - - - (iv)$$

Substitute into (ii):

$$3\pi_0 - 3\pi_1 + 2(5\pi_0 - 2\pi_1) = 0$$

$$\Rightarrow 13\pi_0 - 7\pi_1 = 0 \quad \Rightarrow \pi_1 = \frac{13}{7}\pi_0.$$

Using in (iv)

$$\pi_2 = 5\pi_0 - 2 \cdot \frac{13}{7}\pi_0 = \frac{9}{7}\pi_0.$$

Normalize $\sum \pi_i = 1$

$$\pi_0 + \pi_1 + \pi_2 = 1$$

$$\Rightarrow \pi_0 + \frac{13}{7}\pi_0 + \frac{9}{7}\pi_0 = 1$$

$$\Rightarrow \frac{29}{7}\pi_0 = 1 \quad \Rightarrow \pi_0 = \frac{7}{29}$$

Therefore,

$$\pi_1 = \frac{13}{7} \cdot \frac{7}{29} = \frac{13}{29},$$
$$\pi_2 = \frac{9}{7} \cdot \frac{7}{29} = \frac{9}{29}.$$

Thus

$$\pi = [\pi_0 \quad \pi_1 \quad \pi_2] = \begin{bmatrix} \frac{7}{29} & \frac{13}{29} & \frac{9}{29} \end{bmatrix}.$$

Conclusion:

***Q*-matrix** = $Q = \begin{bmatrix} -5 & 3 & 2 \\ 2 & -3 & 1 \\ 1 & 2 & -3 \end{bmatrix}$ and **Stationary distribution** = $\pi = \begin{bmatrix} \frac{7}{29} & \frac{13}{29} & \frac{9}{29} \end{bmatrix}$.

-: Thank you so much :-