

Cornell Math Explorers' Club 2019: Markov Chains

Instructor's notes

Swee Hong Chan, Ryan McDermott
credits to Kelsey Houston-Edwards

May 4th and 18th 2019

1 Day 1: Introduction to Markov chains

1. (Markov chain icebreaker, at most 10 minutes). Start with the following activity:

- Have the participants stand in circle. First, we roll a six-sided die; hand the dice to the person that many to the left and have them share their name and a fact about them. Then they roll the dice and repeat.
- After a few rounds, start asking questions: Is it possible that Suzy will get the dice 10 times before Joey gets it once? How many times, on average, do you think we will have to roll the dice to get to everyone? The answer is of the order $n \log n$ by a coupon-collecting type argument or Matthew's method, but the proof is too advanced for the participants. We can allude to this fact if some participants are interested in knowing that.
- At some point, most participants would have introduced themselves. If it is needed, change the rules so that the rest of the participants can get at least a turn to introduce themselves.
- Later, after they learn the definition of a Markov chain, ask them to write out why this icebreaker (without the rule change) is a Markov chain, what the state space is, and what the transition probabilities are.

The goal for this activity:

- Encourage participants to be comfortable with each other, and to facilitate a smooth introduction to Markov chains.

2. (Introducing Markov chains, at most 25 minutes).

- (a) Cover the first page of the worksheet. Make sure everyone is on board with the first example. Then give them time to answer the questions.
- (b) After the participants finish the questions, review the answers. Then, give them the second page of the worksheet. Solve the questions in the second page in groups.
- (c) In answering the two questions in the worksheet, use the analogy that the sum of the arrows *away* from a lily pad is one whole pizza, delivered in amounts equal to the corresponding probabilities. The sum of the arrows *toward* a lily pad is the amount of pizza received from neighbors, and this need not add up to a whole pizza.

- (d) Then, explain that this frog example is an example of a Markov chain. The ingredients are states (which are the lily pads), and from each state (lily pad), you (the frog) is given a *random* instruction on how to proceed to the next state (lily pad).

The goal for this activity:

- For the participants to understand the ingredients for a Markov chain, namely the state space and transition probabilities.
- For the participants to understand that the transition probability from any given state should sum up to 1, but the transition probability to any given state does not necessarily sum up to 1.

3. **(More examples of Markov chains, at most 25 minutes).**

- Explain to them about random walks in 1 dimension. Then give out the third page of the worksheet. Ask the participants to write down the probability transition function of this Markov chain.
- Explain to them about random walks in 2 dimensions. Ask the participants to draw the state diagram *and* write down the probability transition function of this Markov chain. Use the analogy that the grid corresponds to buildings in Manhattan, with address given by (street number, avenue number) – if they’ve never seen Cartesian coordinates.
- Explain to them about coin flippings. Ask the participants to come up with the states of the Markov chain, draw the state diagram, *and* write down the probability transition function of this Markov chain.

Goal of this activity:

- To make sure that the participants internalize the concept of Markov chains by giving exercises with ascending levels of difficulties.

4. **(Break, 10 minutes).**

5. **(Multi-steps probability transition function, at most 20 minutes).** Give out the fourth page of the worksheet. Give the participants time (around 10 minutes) to compute the answers in groups. The solution is included below for the convenience of the instructors:

- If a frog starts on lily pad A , what are all the paths it can take in two hops? What is the probability taking each of these paths? (Let them calculate.)
 - path $A \rightarrow A \rightarrow A$, prob $1/3 \cdot 1/3 = 1/9$
 - path $A \rightarrow A \rightarrow B$, prob $1/3 \cdot 1/3 = 1/9$
 - path $A \rightarrow A \rightarrow C$, prob $1/3 \cdot 1/3 = 1/9$
 - path $A \rightarrow B \rightarrow A$, prob $1/3 \cdot 1/2 = 1/6$
 - path $A \rightarrow B \rightarrow C$, prob $1/3 \cdot 1/2 = 1/6$
 - path $A \rightarrow C \rightarrow C$, prob $1/3 \cdot 1 = 1/3$
- Add up all these numbers. What do you get? 1
- Using what you just calculated, what’s the probability of going from lily pad A to lily pad A in two hops? Going from A to B in two hops? A to C in two hops?
 - Lily pad A : $1/9 + 1/6 = 5/18$

- Lily pad B : $1/9$
- Lily pad C : $1/9 + 1/6 + 1/3 = 11/18$

After the participants solved the questions, explain that the objects they just computed are examples of multi-steps probability transition function $P^n(x, y)$, and briefly mention (with minimal explanation) that using matrices (whatever those are) we can compute $P^n(x, y)$ a lot easier.

Goal of this activity:

- To introduce the definition of $P^n(x, y)$ and to motivate using matrices to study Markov chains (in the second week).
6. **(Stationary distribution, at most 20 minutes)**. Give out the fifth page of the worksheet. The questions and the answers are included below:

- Come up with a guess of where the frog starting from lily pad A will end up after 100 hops.

Answer: The probability distribution for the location of the frog after 1 hop is (roughly)

Lily pad A has probability 0.33, Lily pad B has probability 0.33, Lily pad C has probability 0.33.

The probability distribution for the location of the frog after 2 hops is (roughly)

Lily pad A has probability 0.28, Lily pad B has probability 0.11, Lily pad C has probability 0.61.

The probability distribution for the location of the frog after 3 hops is (roughly)

Lily pad A has probability 0.15, Lily pad B has probability 0.1, Lily pad C has probability 0.75.

The probability distribution for the location of the frog after 10 hops is (roughly)

Lily pad A has probability 0.005, Lily pad B has probability 0.003, Lily pad C has probability 0.982.

As we increase the number of hops, if the frog ever enters C, then it will never left. Hence the distribution will move closer and closer to

Lily pad A has probability 0, Lily pad B has probability 0, Lily pad C has probability 1.

Hence it is very likely for the frog to end up at lily pad C after 100 hops.

After solving this problem, discuss the first intuitive meaning as the stationary distribution. This is the limit of the transition function $P^n(x, y)$ as n gets huge. Remember to mention that we would have gotten the same stationary distribution if we have started from B or C instead. Also remember to mention that the stationary distribution sums up to 1.

- Come up with a guess of which side your coin will be after 1000 flips. Is it as clear-cut as in the previous question?

Answer: Flip the coin one time. Then the side of the coin that comes up has probability distribution:

Head has probability $1/2$ to appear, Tail has probability $1/2$ to appear.

Flip the coin again. Then the side of the coin that comes up has probability distribution:

Head has probability $1/2$ to appear, Tail has probability $1/2$ to appear.

Let the participants arrive at the conclusion that the side of the coin after any number of flips always has the probability distribution:

Head has probability $1/2$ to appear, Tail has probability $1/2$ to appear.

This implies that the stationary distribution is given by the distribution above. This shows that neither sides have any advantages that allow them to appear more often.

After solving this problem, discuss the second meaning of the stationary distribution. This is the fraction amount of time you spend in that state as your number of hops gets really huge. In the frog example, the frog spent most of the time at lily pad C (as it will never left that spot), so all the weight of the stationary distribution is focused at C. In the coin example, the coin is at the Head state half of the time, and is at the Tail state half of the time, so the corresponding stationary distribution puts equal weight to state Head and state Tail.

Remember to mention that there is another mathematical definition of the stationary distribution using matrices. This definition makes computing the distribution much simpler, and we will talk more about them next week.

Goal of this activity:

- To introduce two intuitive definitions of the (unique) stationary distribution. Note that we will not mention ergodicity in the class, but allude to it in a non-technical way if this question is being raised. Note that the definition introduced during the second example will be important in solving the knight's move activity (coming right after).

7. (Knight's move activity, at most 25 minutes).

- Give out the sixth page of the worksheet after setting up the problem: On a standard 8×8 chess board with no pieces on it, place a knight on an arbitrary square. The stochastic knight moves by selecting with uniform probability from its legal chess moves. How would you compute the stationary distribution for this problem?
- Then, after some time passes (5 minutes), suggest for them to write down the number of possible moves that lead the stochastic knight to a given square. Wait for the participants to come out with the idea that the stationary distribution of the Markov chain is proportional to those numbers (This is a property of reversible Markov chains).
An intuitive (but non-rigorous!) explanation on why the stationary distribution is proportional to the degree of the vertices: The stationary distribution represents the fraction of time spend on a given vertex. So, if there are only 2 ways a knight can access the corner vertex compared to 8 ways it can access a central vertex, then in the long run the knight will spend 4 times more time at the central vertex than in the corner.
- If there is enough time, tell the participant the fact that the average number of moves it takes for a knight to start at a spot x and return to spot x is equal to $\frac{1}{\pi(x)}$, where π is the stationary distribution of the Markov chain (This is a special property of recurrent Markov chains). Explain that this is how mathematicians solve problems by turning them into other problems, so our new problem is to solve for the stationary distribution.

An intuitive (but non-rigorous!) explanation on why this fact is true: Suppose that the knight starts at the corner, and you make 168 moves. Then, on average, only one of those moves is a return to the corner vertex (by the understanding of stationary distribution as the time spent on any given vertex). Therefore, the average time it takes to return to the corner vertex is 168 moves.

Goal of the activity:

- For the participants to try to solve a question that requires the knowledge of everything they have learned so far about Markov chain.

2 Day 2: Betting on Markov chains

- **(Time filler before the event starts, at most 15 minutes).** Present the following problem: “How many people do we need in this room so that there is at least 50 percent chance of having two people with the same birthday?”. The probability that n people in the room all having a different birthday is

$$p_n = 1 \times \left(1 - \frac{1}{365}\right) \times \left(1 - \frac{2}{365}\right) \times \dots \times \left(1 - \frac{n-1}{365}\right).$$

The first few values of p_n are given by:

$$p_1 = 0\%, \quad p_2 = 99.7\%, \quad p_5 = 97.3\%, \quad p_{10} = 88.3\%, \quad p_{20} = 68.9\%, \quad p_{23} = 49.3\%.$$

Therefore it suffices to have 23 people in the room to guarantee the success probability to be greater than 50 percent. (Print out a figure of the function p_n to show the participants).

1. **(Markov chain icebreaker, 10:00 - 10:10).** Start with the following activity:

- Have the participants write their name on a small piece of paper, and put it inside a box. (Remember to bring a box.) Pick one of the names at random, and have the participant share their name and a fact about them. Then return the name to the box and repeat.
- After a few rounds, start asking questions: Is it possible that Suzy’s name is picked 10 times before Joey gets it once? How many times, on average, will we need to draw the name from the box until everyone introduces themselves? This is the coupon-collector problem, and the answer is $n \log n$. The proof is too advanced for the participants, and we will allude to this fact only if some participants are interested. (NOTE: we have touched on the coupon-collecting problem during Day 1, so some participants might be able to answer this question.)
- After some time, change the rule so that a name is removed if it has been picked before. This is so that each participant has at least one chance to introduce themselves.
- Later, after the participants learn the definition of a Markov chain, mention that the first icebreaker game has Markov property, as the future is influenced by neither the past or the present. On the other hand, the second icebreaker game does not have Markov property, as the future is influenced by the past.

The goals for this activity:

- Encourage participants to be comfortable with each other, especially because some participants were not present on the first day.
- Introducing two models of sampling, one with replacement and one without replacement. We will use these two models as examples for Markov property later in the day.

2. **(Reintroducing Markov chain, 10:10 - 10:30).**

- Cover the first page of the worksheet. Make sure that everyone is on board with the frog-and-coin-toss example.
- Introduce the definition of Markov chain with state space and the probability transition function. Explain that an important property of Markov chain is Markov property (i.e., the future depends only on the present but not the past). Also explain that the first icebreaker game (sampling with replacement) has this property, but not the second icebreaker game.

The goals for this activity:

- For the students to understand the ingredients of Markov chains and the role of Markov property.
- Note that some of the participants did not attend the first day, and hence a detailed re-introduction of Markov chains is necessary here.

3. (**Gambler's ruin, 10:30 - 11:15**). Cover the second and the third page of the worksheet. The answer to the worksheet is included below.

Consider the frog from before that bets on coin tosses. Its betting strategy is to bet one dollar at a time until it makes N dollars (win) or until it runs out of money (lose). Let p_k be the probability that the frog wins the game if it starts with k dollars in its purse.

- What are *all* the ways that the frog can end up winning if it starts with 1 dollar? What is p_1 ?

Answer: There is only one way for the frog to win the game, and the probability is equal to $1/2$. This serves as a warm-up question.

- Suppose that $N = 3$. What are *all* the ways that the frog can end up winning if it starts with 1 dollar? What do these ways have in common?

Answer: The winning moves are right-right, right-left-right-right, right-left-right-left-right-right, etc. The goal is for the participants to realize that the frog keeps coming back to the \$1 lily pad until it decides to move right twice consecutively to win the game. This observation is crucial in solving the next questions.

- Suppose that $N = 3$. Find the exact value of p_1 .

Answer: The goal is for the participants to calculate the probability for every winning move from the previous question, and use that to conclude that

$$p_1 = \frac{1}{4} + \frac{1}{16} + \dots + \frac{1}{4^n} + \dots$$

Use the equilateral triangles interpretation of the geometric series to conclude that $p_1 = \frac{1}{3}$ (bring an illustration of this formula). After this, explain that there is another way to derive p_1 without using infinite series, namely the formula

$$p_1 = \frac{1}{2} \times \left(\frac{1}{2} + \frac{1}{2} p_1 \right).$$

The formula above can be derived using the following explanation: Notice that, the frog must move right for the first step if it wants to win; this contributes to the factor of $\frac{1}{2}$ at the very beginning. For the second step, the frog can move one step right to win the game, OR take a step back to the left and start over again; this contributes to the factor $\left(\frac{1}{2} + \frac{1}{2} p_1 \right)$.

- Suppose that $N = 3$. Give an equation that relates p_1 and p_2 , and use it to find the exact value of p_2 .

Answer: Suppose that the frog starts at the \$2 lily pad. After one step, it either moves one step to the right (which wins the game), or it moves one step to the left and now needs to find a way to win starting from \$1 lily pad. This explanation can be translated into the formula

$$p_2 = \frac{1}{2} + \frac{1}{2} p_1.$$

A similar reasoning for the frog starting at the \$1 lily pad translates into the formula

$$p_1 = \frac{1}{2}p_2.$$

- Suppose that $N = 4$. Find equations that relate p_1 , p_2 , and p_3 . Then find their exact values.

Answer: Use the reasoning from the previous question to derive the equations:

$$\begin{aligned}p_1 &= \frac{1}{2}p_2; \\p_2 &= \frac{1}{2}p_1 + \frac{1}{2}p_3; \\p_3 &= \frac{1}{2}p_2 + \frac{1}{2}.\end{aligned}$$

Let the participants work out that $p_1 = \frac{1}{4}$, $p_2 = \frac{2}{4}$, and $p_3 = \frac{3}{4}$.

- Can you guess the exact value of p_k for any given N ? What equations would you use to show that your guess is correct?

Answer: The value of p_k is given by $p_k = k/N$. This answer is consistent with the equations that we derive from previous questions, namely

$$p_k = \frac{1}{2}p_{k-1} + \frac{1}{2}p_{k+1} \quad (1 \leq k \leq n-1).$$

After this, provide an intuitive explanation on why we expect $p_k = k/N$ is the right answer: Suppose that $k = 1$ and $N = 4$. Note that the frog wins if it ever reaches \$4 lily pad, and loses if it ever reaches \$0 lily pad. Since the \$4 lily pad is three times as far from the starting location compared to the \$0 lily pad, it then makes sense to expect that ending up at the \$0 lily pad (losing) is three times as likely as ending up at the \$4 lily pad (winning). Since the winning probability and the losing probability should add up to 1, it then follows that the winning probability is $1/4$, which is one-third as likely as the losing probability ($3/4$).

Goal of this activity: For the participants to understand how to use geometric arguments and recursions to solve mathematics problems.

4. **(Break, 11:15 - 11:25).**

5. **(Martingale betting system, 11:25 - 11:55).** Give out the third part of the worksheet, and explain the Martingale betting system:

- For the first coin toss, it bets one dollar. That is, the frog gains one dollar if it wins this coin toss, and loses one dollar if it loses this coin toss.
- For the second coin toss, it bets two dollars.
- For the third coin toss, it bets four dollars.
- For the fourth coin toss, it bets eight dollars.
- And so on. Basically, the frog keeps doubling its previous bet until it wins; then it stops. It also stops playing when it doesn't have enough money to play its next bet.

After that, ask the participants to answer the questions in the worksheet. The answers to the worksheet are included below:

- Suppose that the frog starts with the fortune of 10 dollars. What is the fortune of the frog if it wins on the first toss? second toss? third toss? fourth toss?

Answer: The fortune on the event that the first win is in the first three tosses is 11 dollars (10 dollars starting fortune and 1 dollar profit). On the other hand, if the frog lost the first three tosses, then it is left with 3 dollars before the fourth toss, and is forced to drop out from the game.

- Suppose that the frog starts with the fortune of 10 dollars. How would the frog end up losing money in this game? What is the probability that this happens? What if it starts with 20 dollars? with 50 dollars?

Answer: With starting fortune of 10 dollars, the frog needs to lose the first three tosses to end up losing money in the game, and hence the losing probability is $1/8$. With starting fortune of 20 dollars, the frog needs to lose the first four tosses to lose money in the game, and hence the losing probability is $1/16$. With starting fortune of 50 dollars, the frog needs to lose the first five tosses to lose money in the game, and hence the losing probability is $1/32$.

Remind the participants after they answer this question that, as your starting fortune grows larger, your losing probability grows smaller. Also explain to them that this property was one of the driving reason that made this betting strategy popular in the 18th century France, as the gamblers thought (falsely) that this strategy can let them profit in the long run.

- Suppose that the frog visits the casino to play the same game every day with starting fortune 10 dollars. How much money does the frog make if it wins a game? How much money does the frog lose if it loses a game? How much profit (or loss) would it make on average?

Answer: With the starting fortune of 10 dollars, the frog gains 1 dollar most of the time (with probability $7/8$), but loses 7 dollars sometimes (with probability $1/8$). This shows that, the frog does not make or lose money in the long run.

- Redo the previous problem, but with the starting fortune of 20 dollars. Then with the starting fortune of 50 dollars. What do you notice?

Answer: The outcome is that, the frog does make or lose money in the long run with this betting strategy, regardless of the starting fortune.

After the participants answer all the questions in this worksheet, mention that it can be proved (optional stopping theorem) rigorously that, in a fair Markov chain game, the best one can do is to break even in the long run no matter the gambling strategy employed. Then mention that, in an unfair Markov chain game such as the roulette game in the casino (print out a colored picture of a roulette wheel), then one will always lose money in the long run regardless of the strategy employed. Finally, mention that there are non-Markovian games such as blackjack for which this principle does not apply and one can theoretically gain an edge (in this case, by applying card-counting strategy). Also explain that blackjack is non-Markovian as it is an example of sampling without replacement (remind them of the non-replacement icebreaker activity).

Goals of this activity:

- For the participants to understand how to compute expectation.
- For the participants to identify the fallacy behind the claim that the martingale betting system allows one to profit in the long run.

- For the participants to learn that there exists no gambling strategy that can give one an edge in a fair game such as coin toss.

6. (**Monopoly boardgame, 11:55 - 12:15**). Give out the last part of the worksheet and bring out the Monopoly boardgame. (Remember to bring the boardgame.) Ask the participants why this game is not a Markov chain. Some possible answers:

- One reason is because one can get into jail by rolling double this turn if one has rolled double twice before, so the future depends not only on the present but also on the past.
- Another reason is because the cards in the Community Chest are not reshuffled until every card has been used up, which makes it a sampling without replacement and hence not a Markov chain.

After that, show the participants the simulation of the visit probability of the player at any number of steps from the following link:

<http://www.bewersdorff-online.de/amonopoly/> .

(Set up a projector to show the simulation.) Ask the participants why it is much more likely (around five times more likely) for the player to be in Jail than in any other spot. Some factors that contribute to this fact:

- There are rules in the game that force the players to visit jail, such as Community Chest card or landing at “Go to jail” spot.
- There is the rule that one needs to roll a double to escape from the jail (which has success probability $1/6$), which makes it very likely for a player to stuck at the jail for a long time.

Finally, let the participants discuss among themselves how to use what they have learned to design a strategy to play Monopoly. Some possible answers:

- Invest on places with higher visit probability, e.g., it is 45% more likely to land at Illinois Avenue than Park Place, so given that a player own both properties, it will be a better idea to invest on major buildings on the former. Similarly, places that are even steps from the jail are slightly more likely to be visited than the places that are odd steps from the jail. (This is because when one escapes from the jail, one necessarily rolls double.)
- Since the Community Chest cards are (usually) drawn without reshuffling, it is possible to employ card-counting strategy to predict the probability of certain cards being drawn at the later stage of the game.

Goal of this activity:

- To end with a fun example that the participants can discuss among themselves during lunchtime.