

POINTS OF SIGNIFICANCE

Markov models—Markov chains

You can look back there to explain things, but the explanation disappears. You'll never find it there. Things are not explained by the past. They're explained by what happens now.

—Alan Watts

Jasleen K. Grewal, Martin Krzywinski and Naomi Altman

To model biological systems that undergo change, it is not strictly necessary to know the details of the underlying mechanisms. Instead, we can model change as a series of transitions between states. Each transition is assigned a probability that defines the chance of the system changing from one state to another. Together with the states, these transition probabilities define a stochastic model with the Markov property: transition probabilities depend on only the current state—the future can be considered independently of the past if the present is known.

The simplest model with the Markov property is a Markov chain. Consider a single cell that can transition among three states: growth (G), mitosis (M) and arrest (A). At any given time, the cell's state can be represented by a random variable X , which has a value of G, M or A with probability p_G , p_M or p_A , respectively. A Markov chain of this system is a sequence (X_0, X_1, X_2, \dots) , where X_i is the vector of probabilities of finding the system in each state at time step i , and the probability of transitioning from X_i to X_{i+1} depends only on the observed value (G, M or A) of X_i . A realization (observation) of the chain is the set of sequentially observed states (for example, G, M, G, G, ...).

Let's start with a two-state (G, M) Markov chain, which will be discrete-time (time steps are equal) and time-homogeneous (transition probabilities are fixed). At any given time step, the cell in G can undergo mitosis (G to M) with a probability $p_{GM} = 0.2$ or remain in G ($p_{GG} = 1 - 0.2 = 0.8$). We set $p_{GG} > 0$ to allow the cell to stay in G for multiple steps. We also set $p_{MG} = 1$ so that mitosis will take a single time step—once in M, the cell always returns to G in the next step (Fig. 1a). These probabilities define a 2×2 transition probability matrix, T , whose element T_{ij} (i th row, j th column) is the chance of moving from state i to state j . Because the values in the matrix are probabilities, the sum of each row (the probability of being in a given state) must be 1.




Fig. 1 | State transition models, transition matrices T , and the number of transitions required to approximate the steady-state limiting distributions, T^n ($n \rightarrow \infty$), to the displayed number of decimal places. **a, A two-state model in which a cell in the growth phase (G) can undergo mitosis (M) with a probability $p_{GM} = 0.2$. **b**, A three-state model in which the cell may enter temporary arrest (A) from M with a probability of $p_{MA} = 0.2$ but will return to G with a probability $p_{AG} = 0.1$. **c**, A three-state absorption model in which the cell remains in arrest forever. The number of time steps spent in a state before absorption is given by the fundamental matrix, $F_n = (I - Q)^{-1}$, where Q is the highlighted submatrix.**

If our cell starts in G, the initial state probability vector is $X_0 = [p_G = 1, p_M = 0]$. We can obtain subsequent probabilities of states by multiplying the current state probability vector by the transition probability matrix: the second state vector is $X_1 = X_0 T = [0.8, 0.2]$, and the third is $X_2 = X_1 T = X_0 T^2 = [0.84, 0.16]$. Each transition is a multiplication by T , and we obtain the system after n transitions by multiplying by T^n . The elements $(T^n)_{ij}$ are the probabilities of starting in state i and ending in state j after n transitions.

To see what happens in the long run, we calculate T^n in the limit $n \rightarrow \infty$. For our two-state system, T^n converges very quickly (this rate depends on T), and, to two decimal places, T^4 (Fig. 1a) is a good approximation of long-term behavior. Each row i of this matrix gives the probabilities of being in each state after infinitely many time steps, having started in i . These probabilities are the 'limit' of each state and together are called the limiting distribution. Our system always converges to $[0.83, 0.17]$, which tells us that, after a long time, we have a 17% chance of finding the cell in M. This is the case regardless of how we define our initial state, as both rows in T^4 are the

same. Here, the limiting distribution also defines the steady-state behavior of the chain, known as its stationary distribution. Once the stationary distribution is reached, it is unaltered by further transitions.

For irreducible and aperiodic Markov chains, limiting distributions and stationary distributions have unique solutions. They are also equivalent to each other in such chains. In these cases, like in our two-state example, the limiting distribution is independent of the initial state, so the rows of T^n are identical. In an irreducible chain, all states 'communicate' with one another—a given state can be reached from any other state via zero or more transitions. In contrast, an aperiodic chain does not return to a given state in a fixed number of transitions—it may take any number of steps to go from, for example, G to M and then back to G.

Periodic chains may have a stationary distribution but lack a limiting distribution. Such a chain remains in the stationary distribution if it starts there, but otherwise it may never reach it. For example, with $p_{GM} = p_{MG} = 1$, a cell that starts in G will deterministically flip states with each step. A stationary distribution exists— $[0.5, 0.5]$ —but T^n does not converge. In reducible




Fig. 2 | Effect of the initial state (G, M) on the state evolution of 5,000 Markov chains with 20% and 40% chances of arrest. a, The top stacked bar plot shows the state distribution across 60 time steps for 2,500 chains that start in G. The first set of horizontal strips shows arbitrary chains selected from chains with length in the 5th percentile (shortest) to 95th percentile (longest). The vertical black line shows the simulated average absorption time, $n_{G^A} = 30.9$. Results for 2,500 chains that initialize in M are shown below, offset by the average number of steps the G chains initially spend in G. The absorption time for M chains differs by this offset ($25.7 \approx 30.9 - 5$), and their state distribution after the first step (region outlined by a black rectangle) is identical to that for the chains that started in G. **b,** Same as **a**, but with a 40% chance of arrest.

chains, limiting distributions (when they exist) are dependent on the initial state. For example, when $p_{GG} = p_{MM} = 1$, there are two stationary distributions, $[1, 0]$ and $[0, 1]$, for the G and M initial states, respectively.

Now let's extend our model and add the arrest state A, which the cell can enter from M with $p_{MA} = 0.2$ (otherwise it goes back to G). The arrest will be temporary—the cell can return from A to G with a low probability ($p_{AG} = 0.1$) (Fig. 1b). Because this chain is also irreducible and aperiodic, T^n converges, though more slowly than in the previous example. Now the limiting distribution is $[0.625, 0.125, 0.25]$ and the chance of finding the cell in M after a long time is 12.5%.

We can make our chain reducible by making the arrest state permanent, for example, with a drug that inhibits mitosis (Fig. 1c). Now, because all transition probabilities from A are 0 ($p_{AG} = p_{AM} = 0$), we cannot reach G or M from A. Arrest is now an absorbing state, and the chain is absorbing, reducible and aperiodic. The longer this system evolves, the higher the chance of finding the cell in the absorbing state and the lower the probability of finding it in G or M. Regardless of the initial state, $p_G < 0.1$ and $p_M < 0.1$ after 63 transitions, and after 153 transitions the system has converged to within two decimal places of the limiting distribution $[0, 0, 1]$ (Fig. 1c).

Does the rate at which the cell enters arrest depend on the initial state? Does

the cell spend more time in G or M before it goes into arrest? These questions are addressed by the fundamental matrix, $F_n = (I - Q)^{-1}$, where Q is a sub-matrix of T representing the transition probabilities within nonabsorbing states (Fig. 1c), and I is the identity matrix. Summing the values along each row gives us the expected number of time steps before absorption. If we start in G, we expect absorption on average after 30 steps. Of these, 25 will be spent in G and 5 in M. If we start in M, the absorption time decreases to 25 steps (20 in G and 5 in M).

The time to absorption from G is five steps longer than that from M because (a) the average length of time initially spent in G is five steps and (b) once the chain transitions to M, it is statistically identical to the chain that initially started in M, because future states are independent of past states. To explain the value of 5, consider a biased coin with the probability of obtaining a head $p_H = 0.2$. If we treat H as the absorbing state, the expected number of flips to the first head is the absorption time, calculated as the average of the corresponding geometric distribution, $1/p_H = 5$. Moreover, for this coin, $Q = p_T = 0.8$ and $F_n = (1 - Q)^{-1} = 5$, which motivates the matrix form for Q .

To see this explicitly, let's follow the evolution of 2,500 absorbing chains starting in G, with $p_{MA} = 0.2$ (Fig. 2a). Over time, the fraction of cells in A steadily increases,

with a simulated mean time to arrest of 30.9 time steps, close to the expected 30 from **Fn**. Consider now 2,500 chains with the same T but with initial state M. Because $p_{MA} = 0.2$, we expect 20% of those chains to arrest in the second step and 80% to transition to G. Furthermore, the Markov property tells us that the 80% of chains in G will evolve statistically identically to those simulated as starting in G. We show this in Fig. 2a by offsetting the position of the chains that start in M by 5. In fact, if we prepend these chains with the length of time that our cell initially spent in G in the first simulation, we re-create chains with the same statistical properties as those that started in G.

If we increase the absorption probability to $p_{MA} = 0.4$, the values in our fundamental matrix will be halved, and so will the average time to arrest: 15 if we start in G, and 10 if in M (Fig. 2b). As before, the chains that successfully return to G from the initial state M (60%) will be statistically identical from that point to those that started in G.

Markov chains can be generalized to cases of short-term dependency, taking into account recent past states in the chain. For example, the third base in a codon can be probabilistically predicted on the basis of the first two bases by a second-order Markov chain, which 'remembers' the last two states¹.

Discrete-time Markov chains are intuitive, are easy to interpret and have real-life applications in various areas, including weather prediction, finance and biology. They can capture dependencies within a system and reveal interesting long-term behavior. Markov models in which states are not directly observable are called hidden Markov models and will be the subject of our next column. □

Jasleen K. Grewal¹, Martin Krzywinski^{2*} and Naomi Altman³

¹PhD candidate, Canada's Michael Smith Genome Sciences Centre, Vancouver, BC, Canada.

²Staff scientist, Canada's Michael Smith Genome Sciences Centre, Vancouver, BC, Canada.

³Professor of Statistics, The Pennsylvania State University, University Park, PA, USA.

*e-mail: martink@bcgsc.ca

Published online: 30 July 2019
<https://doi.org/10.1038/s41592-019-0476-x>

References

1. Skewes, A. D. & Welch, R. D. *PeerJ* **1**, e127 (2013).

Competing interests

The authors declare no competing interests.