

Topic 5a: Continuous Time/Discrete-Space Stochastic Processes

Learning Objectives

1. Use a transition rate matrix to describe a continuous-time discrete-space Markov Chain (CTDS-MC)
 - a. Describe the properties of the transition rate matrix
 - b. Derive transition probabilities from a transition rate matrix
2. Propose and justify a CTDS-MC for a biological process
 1. Describe 2 different models of molecular evolution.
 2. Describe a phylogenetic tree and the derivation of the tree likelihood under JC69
3. Analyze a CTDS-MC and use these analyses to draw biological conclusions using:
 - a. Simulate a CTDS MC
 - b. Derive and analyze the stationary distribution of a CTDS-MC
 - c. Derive and analyze the Master equations of a CTDS-MC
 - d. Derive and analyze an Ensemble Moment Approximation of a CTDS-MC
 - e. Derive and analyze the diffusion approximation to a CTDS-MC
4. Define reversibility of a CTDS-MC, assess whether a given CTDS is reversible, and give an example of a biological process that is modelled as reversible and why.

Lecture 5.1: Continuous-Time Discrete-Space Stochastic Processes

Review Discrete-Time Discrete-Space Stochastic Processes

Discussion: What are some examples of DTDS stochastic processes?

1. We can draw a diagram of DTDS stochastic processes
2. We can represent DTDS processes with a **transition probability matrix**
3. To analyze DTDS processes we can use: **stationary distributions, hitting times and probabilities, simulation**

Example: 5.1 Species coexistence (revisted)


In ecology, one common application of Markov chains is in modelling the (co)occurrence of species in a community. Let's consider a simple example of a Markov chain modelling the coexistence of two hypothetical species Plant A and Plant B in an ecological reserve.

State 1: Plant A only

State 2: Plant A and B

State 3: Plant B only

1. Draw this stochastic process assuming only one or zero species can be lost/gained in a given year.



2. Let's assume the following transition rate matrix

$$P = \begin{bmatrix} 0.7 & 0.3 & 0 \\ 0.4 & 0.5 & 0.1 \\ 0 & 0.3 & 0.7 \end{bmatrix}$$


What is P_{12} and what does this probability represent.

$P_{12} = 0.3$ is the probability that a Plant B is introduced to an ecosystem and coexists with Plant A that is already present until the next year/time step.

DTDS stochastic processes work well for systems in which are naturally described in discrete time (e.g., annual life cycles) or are only observed by researchers on discrete time intervals (e.g., biannual surveys), but many other systems are better described by events occurring continuously and not synchronously. For example, consider the ecological dynamics described above but in a tropical environments where plants may exhibit reduced synchrony.

An example CTDS process

Let's reimagine the coexistence problem but this time assuming that the plants colonize and are excluded in continuous time. Let c_X be the colonization rate of species X and e_X be the rate at which species X is excluded from an environment containing the other species. We could logically model the system them as:



I have represented these **rates** with grey arrows.

Note there is no longer a arrow of going from State X to State X as this is the rate at which nothing happens which is the default.

Transition Rate Matrices

As we did for the DTDS stochastic process we can describe a CTDS process with a matrix known as the **transition rate matrix**, \mathbf{Q} . In this matrix element q_{ij} $i \neq j$ represents the rate of going from state i to state j whereas element $q_{ii} = -\sum_{j \neq i} q_{ij}$,

Properties of the Transition-Rate Matrix

- The off-diagonal elements of \mathbf{Q} are positive
- The diagonal elements of \mathbf{Q} are negative
- The rows of \mathbf{Q} sum to 0 (Not 1 like the transition probability matrix!)

Example: 5.2 Species coexistence in continuous time

1. Give the CT coexistence model described above, what is the transition rate matrix for this system?

$$\mathbf{Q} = \begin{bmatrix} -c_B & c_B & 0 \\ e_B & -(e_B + e_A) & e_A \\ 0 & c_A & -c_A \end{bmatrix}$$

Deriving the Transition Probability Matrix

Given that the system is in state i at time t , what is the probability that it is in state j at time $t + \Delta t$? In other words, what is:

$$\Pr(X_{t+\Delta t} = j | X_t = i)$$

Note that when $\Delta t = 1$, this is equivalent to calculating the corresponding transition probability.

Recall that *transition probabilities* satisfy the Chapman-Kolmogorob equation:

$$p_{ij}^{t+h} = \sum_k p_{ik}^t p_{kj}^h$$

Subtracting p_{ij}^t from both sides we have:

$$\begin{aligned} p_{ij}^{t+h} - p_{ij}^t &= \left(\sum_k p_{ik}^h p_{kj}^t \right) - p_{ij}^t \\ &= \left(\sum_{k \neq i} p_{ik}^h p_{kj}^t \right) + p_{ii}^h p_{ij}^t - p_{ij}^t \\ &= \left(\sum_{k \neq i} p_{ik}^h p_{kj}^t \right) + (p_{ii}^h - 1) p_{ij}^t \end{aligned}$$

Dividing through by h and taking the limit we have:

$$\lim_{h \rightarrow 0} \frac{p_{ij}^{t+h} - p_{ij}^t}{h} = \left(\sum_{k \neq i} p_{kj}^t \lim_{h \rightarrow 0} \frac{p_{ik}^h}{h} \right) + p_{ij}^t \lim_{h \rightarrow 0} \frac{(p_{ii}^h - 1)}{h}$$

We then note that:

$$\begin{aligned} \lim_{h \rightarrow 0} \frac{p_{ij}^h}{h} &= q_{ij} \\ \lim_{h \rightarrow 0} \frac{p_{ii}^h - 1}{h} &= q_{ii} \\ \lim_{h \rightarrow 0} \frac{p_{ij}^{t+h} - p_{ij}^t}{h} &= \frac{dp_{ij}}{dt} \end{aligned}$$

So:

$$\frac{dp_{ij}}{dt} = \sum_k p_{kj}^t q_{ik}$$

This means we the transition probabilities are given by the system of ODEs represented by the matrix equation:

$$P'(t) = P(t)Q$$

If we ignore for a second that P and Q are matrices, we can recognize that this is a simple linear ODE that we know the solution of (analogous to exponential population growth) and hence we have:

$$P(t) = P(0)e^{Qt}$$

and hence:

$$\Pr(X_{t+\Delta t} = j | X_t = i) = P(t)e^{Q\Delta t}$$

But what does it mean to have e to a matrix? Recall that $e^x = \sum_{n=0}^{\infty} \frac{x^n}{n!}$ so we can express e^{Qt} as:

$$e^{Q\Delta t} = \sum_{n=0}^{\infty} \frac{(Q\Delta t)^n}{n!} \approx \underbrace{\mathbf{I}}_{\text{nothing happens}} + \underbrace{Q\Delta t}_{\text{1 event}} + \underbrace{\frac{Q^2 \Delta t^2}{2}}_{\text{2 events}} + \mathcal{O}(\Delta t^3)$$

where the approximation describes the outcome over a short amount of time Δt . This can be one way of thinking about why the exponential function shows up all the time in biology as it captures and weights the probability of different numbers of events occurring.

Example: 5.3 Species coexistence in continuous time cont

Using the transition rate matrix above:

$$Q = \begin{bmatrix} -c_B & c_B & 0 \\ e_B & -(e_B + e_A) & e_A \\ 0 & c_A & -c_A \end{bmatrix}$$

and letting

$$c_B = 0.3 \frac{1}{\text{year}} \quad c_A = 0.1 \frac{1}{\text{year}} \quad e_B = 0.25 \frac{1}{\text{year}} \quad e_A = 0.01 \frac{1}{\text{year}}$$

What is the probability that an ecosystem starting with both species present (state 2) has only species A present (state 1) after 1 month ($\Delta t = \frac{1}{12}$)?

1. Approximate this value assuming that only one event occurs in time Δt

Here we want the (2, 1) element of the matrix:

$$I + Q\Delta t = \begin{bmatrix} 1 - c_B\Delta t & c_B\Delta t & 0 \\ e_B\Delta t & 1 - (e_B + e_A)\Delta t & e_A\Delta t \\ 0 & c_A\Delta t & 1 - c_A\Delta t \end{bmatrix}$$

which is: $e_B\Delta t = 0.25 * 0.083 = 0.0208$

2. Approximate this value assuming that up to two events occurs in time Δt

Here we want the (2, 1) element of the matrix:

$$I + Q\Delta t + \frac{(Q\Delta t)^2}{2}$$

This gives:

$$\begin{bmatrix} 0.97615 & 0.02383 & 0.00002 \\ 0.01986 & 0.97933 & 0.00081 \\ 0.00017 & 0.00808 & 0.99174 \end{bmatrix}$$

And a resulting probability of $P_{2,1}(1/12) = 0.01986$

Discussion: How do these answers compare and what does that tell us about the validity of the approximation?

The Taylor Series expression above is great as a conceptual understanding and the purposes of simulating a CTDS process incrementally. But what if we want to calculate the transition probability over long time spans? For this we use Eigen-decomposition.

Eigendecomposition

Recall that a square (non-singular matrix) can be decomposed into the product:

$$Q = A.D.A^{-1}$$

where **A** is a matrix whose columns are the (left) eigenvectors of **Q**, **D** is a diagonal matrix of eigenvalues, and **A**⁻¹ is the inverse of **A** which results in a matrix whose rows are the right eigenvalues of **Q**.

$$\mathbf{A} = \begin{bmatrix} \vec{u}_1 & \vec{u}_1 & \dots & \vec{u}_n \end{bmatrix} \quad \mathbf{D} = \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \dots & \dots & \dots & 0 \\ 0 & 0 & \dots & \lambda_n \end{bmatrix} \quad \mathbf{A}^{-1} = \begin{bmatrix} \vec{v}_1 \\ \vec{v}_2 \\ \vdots \\ \vec{v}_n \end{bmatrix}$$

So (without a formal derivation) we have:

$$e^{\mathbf{Q}t} = e^{\mathbf{A} \cdot \mathbf{D} \cdot \mathbf{A}^{-1}t} = \mathbf{A} \cdot e^{\mathbf{D}t} \cdot \mathbf{A}^{-1}$$

where:

$$e^{\mathbf{D}t} = \begin{bmatrix} e^{\lambda_1 t} & 0 & \dots & 0 \\ 0 & e^{\lambda_2 t} & \dots & 0 \\ \dots & \dots & \dots & 0 \\ 0 & 0 & \dots & e^{\lambda_n t} \end{bmatrix}$$

This equality follows can be derived from the fact that e^x can be defined as a series of powers. The derivation itself is cool but not really important for us here.

We can re-order the eigenvalues/eigenvectors in any order, so lets assume that $\lambda_1 > \lambda_2 > \dots$. The eigenvalues of a transition rate matrix have a few convenient properties:

1. There is at least one eigenvalue of $\lambda = 0$ if Q is "strongly connected" (it is possible to reach every state from every other state eventually, this is necessary for there to be a unique stationary distribution).
2. All the other eigenvalues are less then 0, $0 > \lambda > 2 \min q_{i,i}$

Then recall that our goal here is to understand changes over large amounts of time, in which case:

$$e^{\lambda_1 t} = 1 \gg e^{\lambda_2 t} \gg \dots$$

So we can approximate:

$$e^{\mathbf{Q}t} \approx \begin{bmatrix} \vec{u}_1 & \vec{u}_1 & \dots & \vec{u}_n \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \dots & \dots & \dots & 0 \\ 0 & 0 & \dots & 0 \end{bmatrix} \cdot \begin{bmatrix} \vec{v}_1 \\ \vec{v}_2 \\ \vdots \\ \vec{v}_n \end{bmatrix} = \underbrace{\vec{u}_1 \cdot \vec{v}_1}_{\approx \mathbf{P}}$$

Note that \vec{u}_1 is a column vector and \vec{v}_1 is a row vector so their product is a matrix. The net result is approximately the transition probability matrix over long periods of time \mathbf{P} . Note that the result here doesn't depend on t and hence is the transition matrix in the realm of time over which the **stationary distribution** applies.

What are \vec{u}_1 and \vec{v}_1 ? \vec{u}_1 is known as the left eigenvector and \vec{v}_1 as the right eigenvector. By definition of an eigenvector we have:

$$\begin{aligned} \mathbf{Q} \cdot \vec{u}_1 &= \lambda \vec{u}_1 = \vec{0} \\ \vec{v}_1 \cdot \mathbf{Q} &= \lambda \vec{v}_1 = \vec{0} \end{aligned}$$

Recall that the leading eigenvalue is 0. This gives two systems of equations:

$$\sum_k u_{1,k} Q_{j,k} = 0 \quad \forall j \quad \sum_j Q_{j,k} v_{1,j} = 0 \quad \forall k$$

Given that the rows of \mathbf{Q} sum to 0 by definition $\sum_k Q_{j,k}$ we know that $u_{1,1} = u_{1,2} = \dots = u_{1,n}$ and the row vector \vec{v}_1 (the leading right eigenvector) is the stationary distribution!

Example: 5.4 Species coexistence in continuous time cont

Python: Lecture5_1.ipynb

1. What is the probability of going from state 2 to state 1 in the long term?

Discussion: How does this answer compare to the results over the short term?

2. What is the long-term probability of ending up in state 1?

Lecture 5.2 Models of Molecular Evolution

One important application of CTDS models in biology is in describing how genomes evolve. Specifically, in describing how mutations occur at different bases in the genome. Such models are known as **models of molecular evolution**. There are two general types of such models: "nucleotide substitution" models and "amino acid substitution" models. Nucleotide models describe the rate at which each of the four nucleotides (A,C,G,T) mutates to become each of the other nucleotides. Amino-acid models describe how each of the 20 essential Amino acids evolve to become each of the other 20.

The motivation and design of AA models is complicated having to do with when and how mutations occur and the relative chemical properties of different AAs. The standard models for AA substitution are called the **BLOSUM** matrices (<https://en.wikipedia.org/wiki/BLOSUM>). This is all I am going to mention about this.

There are a array of Nucleotide substitution models, each with different properties and assumptions. In this class we are going to talk about three different options, each of which is known by an acronym: JC69, HKY84, GTR. Each of these is by definition a 4 state process.

JC69: Jukes Cantor 1969

This is the simplest model of molecular evolution. It assumes that each nucleotide is equally likely to become each other nucleotides. Specifically the rate of going from nucleotide x to nucleotide y is denoted as μ .

Example: 5.5 JC69

1. Draw a transition diagram for the JC69 model.

2. What is the transition rate matrix for JC69?

$$Q^{JC69} = \begin{bmatrix} -3\mu & \mu & \mu & \mu \\ \mu & -3\mu & \mu & \mu \\ \mu & \mu & -3\mu & \mu \\ \mu & \mu & \mu & -3\mu \end{bmatrix}$$

3. Suppose that $\mu = 10^{-3} \frac{\text{mut}}{\text{site} \cdot \text{year}}$ (a value appropriate for the evolution of bases in a virus genome). Suppose it takes approximately 2 weeks ($\Delta t = \frac{1}{24}$) for patient A to infect patient B. What is the probability that a mutation occurred at a focal site during the course of patient A's infection? What about two mutations?

The per-base pair mutation rate is μ meaning that the total rate at which a mutations occur is 3μ (see the off diagonal element of the rate matrix). We need to ask, what is the probability that a mutation occurs at or before $t = \frac{1}{24}$?

The CDF of the exponential distribution is:

$$\Pr(T \leq t^* = \frac{1}{24}) = 1 - e^{-\lambda t^*} = 1.24 \times 10^{-4}$$

Where $\lambda = 3\mu$

The probability that two mutations occurred can be calculated from the CDF of the Erlang distribution with $k = 2$ (this CDF is a function called a Gamma Regularized Function, so we will just evaluate it numerically)

$$\Pr(2 \text{ mutations}) = 7.81 \times 10^{-9}$$

To complement this let's calculate the 1-step and 2-step approximations to the transition probability matrix.

$$P_1(\Delta t = \frac{1}{24}) = \mathbf{I} - \mathbf{Q}\Delta t = \begin{bmatrix} 1 - 3\mu\Delta t & \mu\Delta t & \mu\Delta t & \mu\Delta t \\ \mu\Delta t & 1 - 3\mu\Delta t & \mu\Delta t & \mu\Delta t \\ \mu\Delta t & \mu\Delta t & 1 - 3\mu\Delta t & \mu\Delta t \\ \mu\Delta t & \mu\Delta t & \mu\Delta t & 1 - 3\mu\Delta t \end{bmatrix}$$

$$= \begin{bmatrix} 0.99987 & 0.000042 & 0.000042 & 0.000042 \\ 0.000042 & 0.99987 & 0.000042 & 0.000042 \\ 0.000042 & 0.000042 & 0.99987 & 0.000042 \\ 0.000042 & 0.000042 & 0.000042 & 0.99987 \end{bmatrix}$$

$$P_2(\Delta t) = \begin{bmatrix} 6\Delta t^2\mu^2 - 3\Delta t\mu + 1 & \Delta t\mu(1 - 2\Delta t\mu) & \Delta t\mu(1 - 2\Delta t\mu) & \Delta t\mu(1 - 2\Delta t\mu) \\ \Delta t\mu(1 - 2\Delta t\mu) & 6\Delta t^2\mu^2 - 3\Delta t\mu + 1 & \Delta t\mu(1 - 2\Delta t\mu) & \Delta t\mu(1 - 2\Delta t\mu) \\ \Delta t\mu(1 - 2\Delta t\mu) & \Delta t\mu(1 - 2\Delta t\mu) & 6\Delta t^2\mu^2 - 3\Delta t\mu + 1 & \Delta t\mu(1 - 2\Delta t\mu) \\ \Delta t\mu(1 - 2\Delta t\mu) & \Delta t\mu(1 - 2\Delta t\mu) & \Delta t\mu(1 - 2\Delta t\mu) & 6\Delta t^2\mu^2 - 3\Delta t\mu + 1 \end{bmatrix}$$

4. Of course a virus genome has more than 1 site. SARS-CoV-2 has ≈ 1000 bases. What is the probability that at least one mutation has occurred at any one site during the course of patient A's infection.

Now $\lambda = 3 * 1000 * 10^{-3} = 3$.

The probability that there is at least one mutation then is: $1 - e^{-\frac{3}{24}} = 0.117$


The probability that two mutations occurred is: 0.007.

So the probability that exactly one mutation occurring is: $\approx 0.117 - 0.007 = 0.11$ assuming that three or more mutations are very unlikely to occur.

Discussion: JC69 is an extreme approximation of genome evolution, what would you add/change about this model for realism? Note that the goal is also to minimize parameters that have to be estimated.

K80: Kimura 1980 & HKY84: Hasegawa-Kishino-Yano 1984

JC69 assumes that all mutations are equally likely to occur, this is certainly not the case. We can model some of this complexity by



There are two general types of base pairs "purines" which are large and "pyrimidines" which are small. In order to maintain the stability of the genome mutations are more likely to occur within these groups (purine->purine or pyrimidine->pyrimidine), mutations known as transitions, then between them, transversions. Hence we propose the following model known as K80:

$$\mathbf{Q}^{K80} = \begin{bmatrix} -(\alpha + 2\beta) & \alpha & \beta & \beta \\ \alpha & -(\alpha + 2\beta) & \beta & \beta \\ \beta & \beta & -(\alpha + 2\beta) & \alpha \\ \beta & \beta & \alpha & -(\alpha + 2\beta) \end{bmatrix}$$

Where the states are ordered (1: T, 2: C, 3:A, 4:G) and transitions occur at a rate α and transversions occur at a rate β .

Example: 5.6 K80

1. What is the stationary distribution of the K80 model above?

Recall that the leading eigenvalue (of a strongly connected transition rate matrix like this one) is $\lambda = 0$, so we just want to solve for the corresponding right eigenvectors.

$$\vec{v}_1 \cdot \mathbf{Q} = \lambda \vec{v}_1 = 0$$

$$\begin{aligned}
 \begin{bmatrix} 0 & 0 & 0 & 0 \end{bmatrix} &= \begin{bmatrix} v_1 & v_2 & v_3 & v_4 \end{bmatrix} \cdot \begin{bmatrix} -(\alpha + 2\beta) & \alpha & \beta & \beta \\ \alpha & -(\alpha + 2\beta) & \beta & \beta \\ \beta & \beta & -(\alpha + 2\beta) & \alpha \\ \beta & \beta & \alpha & -(\alpha + 2\beta) \end{bmatrix} \\
 &= \begin{bmatrix} v_1(-\alpha - 2\beta) + \alpha v_2 + \beta v_3 + \beta v_4 \\ v_2(-\alpha - 2\beta) + \alpha v_1 + \beta v_3 + \beta v_4 \\ v_3(-\alpha - 2\beta) + \alpha v_4 + \beta v_1 + \beta v_2 \\ v_4(-\alpha - 2\beta) + \alpha v_3 + \beta v_1 + \beta v_2 \end{bmatrix}^T
 \end{aligned}$$

and hence $v_1 = v_2 = v_3 = v_4$. Recall that eigenvectors have arbitrary length but we want to normalize so that the eigenvector sums to 1, so...

$$\vec{\pi} = \left[\frac{1}{4} \quad \frac{1}{4} \quad \frac{1}{4} \quad \frac{1}{4} \right]$$

Discussion: What does this stationary distribution mean for the frequency of A, C,G,T in the genome in the long term?

2. Consider a sequences with 50 bases, simulate molecular evolution for $t = 20$ time steps assuming $\alpha = 0.1$ and $\beta = 0.04$.

Phython: [Lecture5_2.ipynb](#)

3. Do the observed frequencies of each nucleotide at the end of your simulation match your expectations from the stationary distribution?

Phython: [Lecture5_2.ipynb](#)

This model predicts that the frequencies of A,C,G, and T's observed in the genome in the long term should be equal. Given that evolution has been occurring for a long time, this means that all genomes should be made up of approximately equal amounts of each nucleotide. This is not true though in the real world, and in fact the proportion of each nucleotide varies significantly across the tree of life. Hence, the HKY84 model takes the backbone of K80 and adjusts it so that the stationary distribution reflects an "empirical distribution".

Specifically, let $\vec{\pi}$ be a vector representing the OBSERVED proportion of each nucleotide in the genome. Then the HKY84 model is represented by the transition rate matrix:

$$\mathbf{Q}^{HKY} = \begin{bmatrix} -(\alpha\pi_C + \beta(\pi_A + \pi_G)) & \alpha\pi_C & \beta\pi_A & \beta\pi_G \\ \alpha\pi_T & -(\alpha\pi_T + \beta(\pi_A + \pi_G)) & \beta\pi_A & \beta\pi_G \\ \beta\pi_T & \beta\pi_C & -(\alpha\pi_G + \beta(\pi_T + \pi_C)) & \alpha\pi_G \\ \beta\pi_T & \beta\pi_C & \alpha\pi_A & -(\alpha\pi_A + \beta(\pi_C + \pi_T)) \end{bmatrix}$$

The stationary distribution of this model is:

$$\vec{\pi} = [\pi_T, \pi_C, \pi_A, \pi_G]$$

Or the empirical stationary distribution.

GTR: General Time Reversible

(Model originally developed by Simon Taveré in 1986)

By convention in this model the nucleotides are listed in alphabetical order so the states are in order $\{A, C, G, T\}$

$$Q^{GTR} = \begin{bmatrix} -(a\pi_C + b\pi_G + c\pi_T) & a\pi_C & b\pi_G & c\pi_T \\ a\pi_A & -(a\pi_A + d\pi_G + e\pi_T) & d\pi_G & e\pi_T \\ \pi_A b & \pi_C d & -(\pi_A b + \pi_C d + f\pi_T) & f\pi_T \\ \pi_A c & \pi_C e & f\pi_G & -(\pi_A c + \pi_C e + f\pi_G) \end{bmatrix}$$

A CTDS stochastic process is known as **Time Reversible** if it satisfies:

$$\pi_i Q_{ij} = \pi_j Q_{ji} \quad \forall i, j$$

This implies:

$$\begin{aligned} \pi_i \mathbf{P}_{ij}(t) &= \pi_j \mathbf{P}_{ji}(t) \\ \mathbf{P}_{ij}(t) &= \frac{\pi_j}{\pi_i} \mathbf{P}_{ji}(t) \end{aligned}$$

More intuitively a stochastic process is time reversible if:

$$\begin{aligned} \Pr(X_t = x_0, X_{t+\delta_1} = x_1, X_{t+\delta_1+\delta_2} = x_2, \dots, X_{t+\sum_j^n \delta_j} = x_n) = \\ \Pr(X_t = x_n, X_{t+\delta_n} = x_{n-1}, X_{t+\delta_n+\delta_{n-1}} = x_{n-2}, \dots, X_{t+\sum_j^n \delta_j} = x_0) \end{aligned}$$

Example: 5.7

1. Show that the GTR model is time reversible

Check element $A \leftrightarrow C$:

$$\begin{aligned} \pi_A Q_{A,C} &= \pi_C Q_{C,A} \\ \pi_A a\pi_C &= \pi_C a\pi_A \end{aligned}$$


A similar process could be done for the other bases.

Lecture 5.3 Tree Likelihoods

Phylogenetic Trees

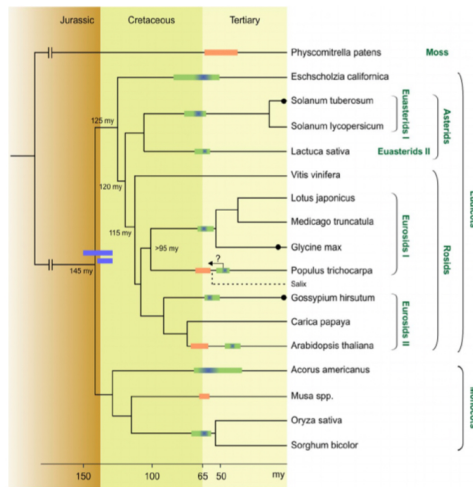
A phylogenetic tree is a diagrammatic representation of the evolutionary relationships among a group **taxa** (distantly related organisms, for example different species, different viruses as viruses do not reproduce sexually and evolve rapidly, or individuals from distinct populations). It depicts the common ancestry and divergence of these taxa over time.

Apes



Yang + Rannala 2005

Flowering Plants



Fawcett et al 2015 PNAS

Discussion: Given this tree, what is the closest relative of a human?

There are multiple different uses of phylogenetic trees:

Classification and Taxonomy: Phylogenetic trees provide a framework for classifying and organizing biodiversity. By identifying evolutionary relationships, biologists can create taxonomies that reflect the true genetic or evolutionary relatedness among species.

Comparative Biology: Phylogenies provide a framework for comparative biology, allowing researchers to study the diversity of traits, behaviours, and physiological processes across related species while accounting for their evolutionary relationships.

Evolutionary Processes: Phylogenies help researchers investigate the mechanisms of evolutionary change, such as the relative strengths of natural selection and genetic drift, or the drivers of speciation and extinction.

Phylogenetics: Inference of the rate at which lineages speciate/go extinct. In the context of viral phylogenetic trees this can inform the transmission rate and recovery rate of an infectious disease.

Here are some key components and concepts related to phylogenetic trees:

1. **Nodes:** Represent points where lineages split, indicating common ancestors. Nodes are often labeled with the estimated time of divergence or other relevant information (such as the statistical certainty of that node).
2. **Branches:** Connect nodes and represent the evolutionary pathways from common ancestors to descendant species or taxa.
3. **Tips or Leaves:** Represent the terminal endpoints of the tree, indicating currently existing taxa.
4. **Root:** The point on the tree that represents the common ancestor of all the taxa included in the tree.
5. **Branch Lengths:** The lengths of branches can represent evolutionary distances, such as genetic divergence or time. Longer branches indicate more evolutionary change (e.g., mutations). In this class we will use branch lengths to indicate time.

Discussion: How many tips and how many nodes are in the ape tree above? Which two species are most closely related?

Phylogenetic trees can be constructed based on various types of data, including molecular sequences (DNA, RNA, or protein), morphological characteristics, or a combination of both. Tree-building methods use algorithms to analyze the data and infer the most likely evolutionary relationships among the taxa. There are three general classes of tree building methods: Parsimony, Likelihood, and Bayesian. Here we will only talk about the second and only in the case of DNA (i.e. molecular) data.

When constructing a phylogenetic tree, we often consider a *single genome per taxa* known as the **consensus sequence**. Note that species have a lot of diversity within them as we explored with the coalescent. But for a phylogeny we ignore this and go with

a single "most representative" sequence.

Calculating the tree likelihood

We will talk much more about this in Topic 7, but here we want to calculate the Likelihood of a model, Θ , given some observed data \mathcal{D} :

$$\mathcal{L}(\Theta|\mathcal{D}) \stackrel{\text{def}}{=} \Pr(\mathcal{D}|\Theta)$$

In our case the data is the observed genome sequences and the model is the parameters of a molecular evolution transition rate matrix.

Example: 5.8: Molecular models

1. Assuming a JC69 model what is Θ ?

$$\Theta^{JC69} = \{\mu\}$$

2. Assuming a HKY84 model what is Θ ?

$$\Theta^{HKY84} = \{\alpha, \beta\}$$


Note that $\vec{\pi}$ is obtained directly from the data and is not a flexible parameter.

Example: 5.9: Tree Likelihood

Consider the following phylogenetic tree.

Newick Format

$$((G : t_1, G : t_1) : t_2, C)$$



Parameters:

$$\begin{aligned}\mu &= 1.2 \\ t_1 &= 0.3 \\ t_2 &= 0.5\end{aligned}$$

1. What is the likelihood of this tree given the observed sequences under the JC69 model?

Let \mathcal{D} be the "Data" which is the states at the tips of the tree.

Let Θ^{JC69} be the parameters of the JC69 model

Let \mathcal{T} be the tree topology (branching pattern) and branch lengths.

The likelihood of the tree given the data is the probability of the data given the tree:

$$\mathcal{L}(\mathcal{T}, \Theta | \mathcal{D}) = \Pr(\mathcal{D} | \mathcal{T}, \Theta)$$

To calculate the probability of the tree we:

- Calculate the probability of the changes on each branch (edges in the tree graph).
- When there is an internal node that we don't know the state of, we simply sum over all the possible values it could take on.
- We weight the root node states by the stationary distribution $\Pr(Y) = \pi_Y$

$$\mathcal{L}(\mathcal{T}, \Theta | \mathcal{D}) = \Pr(\mathcal{D} | \mathcal{T}, \Theta) = \sum_{Y \in \{A, C, G, T\}} \Pr(Y) P_{Y,C}(t_1 + t_2) \sum_{X \in \{A, C, G, T\}} P_{X,Y}(t_2) P_{G,X}(t_1) P_{G,X}(t_1)$$

So how do we calculate $P_{i,j}(t)$? Recall that the transition probability matrix of a CTDS process is given by:

$$\mathbf{P}(t) = e^{\mathbf{Q}t}$$

We can solve for this numerically in python.

[Python: Lecture5_3.ipynb](#)

2. What is the probability that root node Y is a G ?

To calculate this we use a modified version of Bayes Theorem that allows for additional conditioning of all terms:

$$\Pr(A|B, C) = \frac{\Pr(B|A, C) \Pr(A|C)}{\Pr(B|C)}$$

Here we consider a specific tree \mathcal{T} (see the tree topology above) and molecular model Θ (JC69 with $\mu = 1.2$) so:

$$\Pr(Y = G | \mathcal{D}, \mathcal{T}, \Theta) = \frac{\Pr(\mathcal{D} | Y = G, \mathcal{T}, \Theta) \Pr(Y = G | \mathcal{T}, \Theta)}{\Pr(\mathcal{D} | \mathcal{T}, \Theta)}$$

where \mathcal{D} is the 'data' at the tips.

Here $\Pr(Y = G | \mathcal{T}, \Theta) = \pi_G$ is the "prior" probability that the root is a G, assuming that evolution has occurred for a long time prior to this tree our best guess for the state is the stationary distribution of the JC69 model or $\pi_G = \frac{1}{4}$

The denominator $\Pr(\mathcal{D} | \mathcal{T}, \Theta)$ is the likelihood we calculated in the last part.

Finally, $\Pr(\mathcal{D} | Y = G, \mathcal{T}, \Theta)$ is very similar to the likelihood above but fixing the root state at a G

$$\Pr(\mathcal{D} | Y = G, \mathcal{T}, \Theta) = P_{G,C}(t_1 + t_2) \sum_{X \in \{A, C, G, T\}} P_{X,G}(t_2) P_{G,X}(t_1) P_{G,X}(t_1)$$

When there is more than one site/base being considered, we assume that evolution is independent at each base. There are more complex models that allows the mutation rate to vary from base to base, but lets assume it is equal for each base for now.

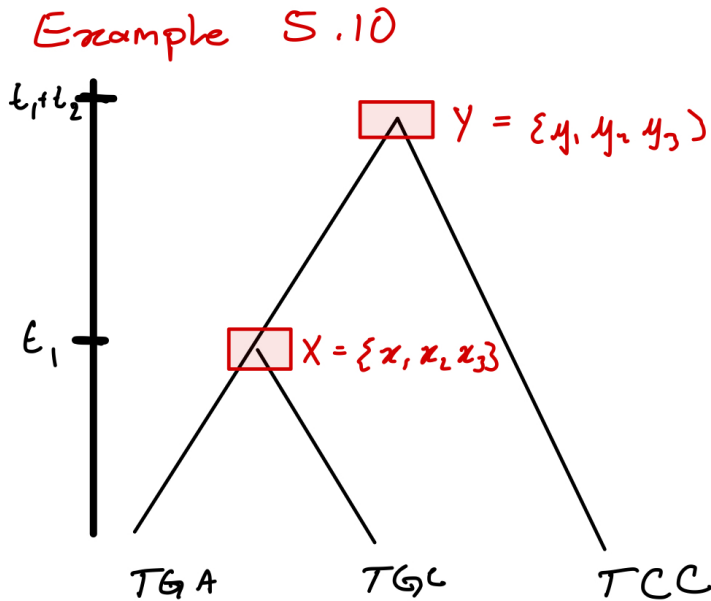
Example: 5.10: Tree Likelihood 2

Python: [Lecture5_3.ipynb](#)

Consider the following phylogenetic tree. What is the likelihood of observing the following phylogenetic tree?

Newick Format

$((TGA : t_1, TGC : t_1) : t_2, TCC)$



When considering multiple sites/bases, we assume that each of them evolves independently such that:

$$\begin{aligned} \Pr(((TGA : t_1, TGC : t_1) : t_2, TCC)) &= \Pr(((T : t_1, T : t_1) : t_2, T)) \\ &\times \Pr(((G : t_1, G : t_1) : t_2, C)) \\ &\times \Pr(((A : t_1, C : t_1) : t_2, C)) \end{aligned}$$

Lecture 5.4 Simulations and Stochastic Mapping

Simulating Evolution on a Tree

Here we discussed how to simulate genome sequences that are consistent with a given **tree topology** (the branching pattern and branch lengths of a tree). But, as in the coalescent, we often know the genome sequences at the present day but want to understand who the ancestors are and particular trajectories of mutations that may have given rise to the OBSERVED sequences.

Stochastic Mapping

- Method from Nielsen 2002

Stochastic mapping is a method for 1) simulating molecular evolution on a tree and 2) inferring the ancestral sequence at internal nodes in the tree.

Example: 5.11: Ancestral Reconstruction

1. Consider the following tree, propose a likely ancestor at node X and node Y? How much does each tip in the tree inform the likely sequence of each ancestor?

Conclusion: Each tip in the tree contains information about each node in the tree proportional to the inverse of the distance between them.

2. Where could mutations have happened in the tree to be consistent with the state of the tips?

Stochastic mapping provides a method for answering the above questions in a manner that is consistent with a given model of molecular evolution. This procedure has four steps:

Step 1: Post-order traversal calculation of node probabilities

- When looking at probabilities and stochastic processes on trees we consider two different types of calculations: **post-order traversal** and **pre-order traversal**.

Pre-order traversal: calculations that proceed from the root through the descendants to the tips.

Post-order traversal: calculations that proceed from the tips through the ancestors to the root.

Let k be an index over all the internal nodes in a tree.

Let D_k be the data (sequences) of all the descendant tips from node k

Let Y_k be the (unobserved) state at internal node k

First e calculate $\Pr(D_k|Y_k = i) = f_{k,i}$

Example: 5.11: Ancestral Reconstruction

3. Calculate $f_{X,i}$ and f_{Y_i} What information is being used to do these calculations?

Step 2: Root Probability

At the root node r calculate the probability

$$\Pr(Y_r = i|D_r = D) = \frac{f_{r,i}\pi_i}{\sum_j f_{r,j}\pi_j}$$

where π_j is the stationary distribution of the molecular evolution model or, if applicable, also equal to the empirical distribution of states.

Note that at the root D_r is all the data so we can also just call this D .

Example: 5.11: Ancestral Reconstruction

4. What is the root node in this case?

5. Calculate $\Pr(Y_r = i|D_r = D)$ for the root node. What information is being used here?

Step 3: Pre-order traversal calculation of node probabilities

Now we travel down the tree calculating the probability that each node is in a given state given the ancestor.

Notationally, for a node $k - 1$ let node k be its direct ancestor.

Let t_{k-1} be the length of the branch between nodes $k - 1$ and k

$$\Pr(Y_{k-1} = j|Y_k = i, D) = \frac{f_{k-1,j}P_{i,j}(t_{k-1})}{\sum_k f_{k-1,k}P_{i,k}(t_{k-1})}$$

Example: 5.11: Ancestral Reconstruction

6. What is/are the internal nodes in this case and who are their direct ancestors.

7. Calculate $\Pr(Y_{k-1} = j|Y_k = i, D) \quad \forall$ internal nodes

Step 4: Simulation of ancestral states and mutations

Step 4A: Given $\Pr(Y_r = i | D_r = D)$ and $\Pr(Y_{k-1} = j | Y_k = i, D)$, simulate the ancestor at each internal node.

Example: 5.11: Ancestral Reconstruction

8. Simulate the state of the internal nodes in the tree.

Step 4B: Given the states at the nodes and the tips, simulate mutational trajectories along the branches that are consistent with the tip states.

Topic 5b: Continuous Time/Discrete-Space Stochastic Processes

Lecture 5.5 Master Equations

Review:

- CTDS stochastic processes can be visualized using a transition diagram
- CTDS stochastic processes can be formulated using a transition rate matrix
- The right-eigenvector, \vec{v}_1 , determines the stationary distribution of the CTDS process
- Models of molecular evolution are useful examples of CTDS processes
- We can use molecular evolution models to formulate the probability a given tree gave rise to the observed sequences (the tree likelihood).
- Stochastic mapping using molecular models to predict ancestral states and simulate mutations on a tree.

So far we have considered only two methods for analyzing a CTDS processes: simulation and stationary distributions. The next few lectures will discuss alternative analysis approaches.

Definition

A **master equation** is a differential equation describing the dynamics of the state probability distribution, $P_n(t)$ for a CTDS stochastic process.

Solving the system of master equations for all states in the process then, gives the **exact solution** for the predicted outcome represents the density of an infinite number of simulation runs.

Let $P_n(t)$ be the probability that the system is in state n at time t .

$$\begin{aligned} \frac{dP_n(t)}{dt} &= -P_n(t) \underbrace{\sum_{i \neq n} q_{n,i}}_{\text{flow out}} + \underbrace{\sum_{i \neq n} P_i(t) q_{i,n}}_{\text{flow in}} \\ &= P_n(t) q_{n,n} + \sum_{i \neq n} P_i(t) q_{i,n} = \sum_i P_i(t) q_{i,n} \end{aligned}$$

Example: 5.12 JC69

1. Formulate the system of master equations describing the dynamics of JC69 model.

The system is particularly easy in this case as the rates are independent of the state.

$$\frac{dP_n}{dt} = -3\mu P_n(t) + \sum_{i \neq n} \mu P_i(t) \quad n = \{1, 2, 3, 4\}$$

2. Assuming that the system starts in state i what is the probability of being in state j after t units of time?

We can fortunately solve this system of equations exactly:

$$\begin{aligned} P_n(t) &= \left(\frac{1}{4} + \frac{3}{4} e^{-4\mu t} \right) P_n(0) + \left(\frac{1}{4} - \frac{1}{4} e^{-4\mu t} \right) \sum_{i \neq n} P_i(0) \\ P_{i,j}(t) &= P_j(t | P_i(0) = 1) = \begin{cases} \left(\frac{1}{4} + \frac{3}{4} e^{-4\mu t} \right) & \text{if } i = j \\ \left(\frac{1}{4} - \frac{1}{4} e^{-4\mu t} \right) & \text{if } i \neq j \end{cases} \end{aligned}$$

Given that this has a known general solution makes it MUCH faster and easier to calculate tree likelihoods under this model as we don't need to do matrix exponentiation to get $P_{i,j}(t)$.

3. After 1 unit of time, how close is the system to the stationary distribution?

Note that as $t \rightarrow \infty$ we have:

$$\hat{P}_n = \left(\frac{1}{4}\right) P_n(0) + \left(\frac{1}{4}\right) \sum_{i \neq n} P_i(0) = \frac{1}{4}$$

which is what we knew to be the stationary distribution. After one unit of time we have:

$$P_n(1) = \left(\frac{1}{4} + \frac{3}{4}e^{-4\mu}\right) p_n(0) + \left(\frac{1}{4} - \frac{1}{4}e^{-4\mu}\right) \sum_{i \neq n} P_i(0)$$

So using $|P_n(1) - \hat{P}_n|$ as a measure of closeness we have:

$$\left|P_n(1) - \hat{P}_n\right| = \frac{3}{4}e^{-4\mu}P_n(0) + \frac{1}{4}e^{-4\mu}P_i(0)$$

Often there is no known general solution and we must solve the dynamics numerically.

Example: 5.13 SIS model

Python: [Lecture5_5.ipynb](#)

Consider an SIS stochastic epidemic where transmissions occur at a mass-action rate of $\frac{\beta}{\kappa} * S * I$ with $\beta = 0.5$, hosts recover (becoming susceptible again) at a rate $\gamma = 0.1$ and the total population size is $\kappa = 100$.


1. Numerically solve the system of master equations for $t = [0, 50]$ assuming that initially there are $I(0) = 5$ infections.

We can describe this process with a 1D state space counting the number of infected hosts between 0 and 100. If the number of infected hosts is given by the state n , the number of susceptible hosts is given by $\kappa - n$.

To derive the master equation note that the system can only transition between neighbouring states through recovery and infection. The system leaves state n via infection (at rate $\frac{\beta}{\kappa}n(\kappa - n)$) going to state $n + 1$ and recovery (at rate γn going to state $n - 1$) and enters state n from state $n - 1$ by infection (at rate $\frac{\beta}{\kappa}(n - 1)(\kappa - n + 1)$) and via recovery from state $n + 1$ (at rate $\gamma(n + 1)$).

The system of master equations is:

$$\frac{dP_n(t)}{dt} = -P_n(t) \left(\gamma n + \frac{\beta}{\kappa} n(\kappa - n) \right) + P_{n-1}(t) \frac{\beta}{\kappa} (n - 1)(\kappa - n + 1) + P_{n+1}(t) \gamma (n + 1)$$



2. Use the master equations to calculate the probability that the disease goes extinct by time t .

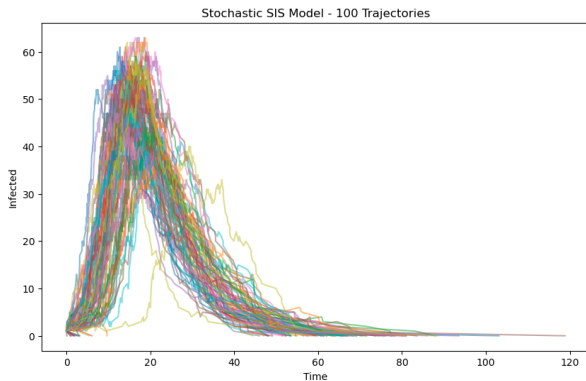
The disease goes extinct if $I(t) = n = 0$.

When $n = 0$ the Master equation reduces to:

$$\frac{dP_0(t)}{dt} = P_1(t)\gamma \quad \text{where } P_0(0) = 0 \text{ and } P_1(0) = 1$$

$$P_0(t) = \int_0^t \gamma P_1(x) dx$$

3. Simulate 100 trajectories from the same stochastic process. How do your results compare to the solution from the master equations? What are some advantages/disadvantages to simulations versus master equations?



In many ways the trajectories are faster, they also give us an idea of the dynamics of individual outcomes in a way that the master equation do not. But they are finite in number and hence do not provide an exact result. Plus we would have to infer from this data the probability of extinction for example. If this probability is small it can be hard to obtain accurately from trajectories.

The state space of many biological stochastic processes may be more conveniently viewed in 2 dimensions rather than a single dimension (note that a 2 D discrete state space can always be mapped into 1D but the mapping might be complicated an inconvenient). As an example of this consider the SIR epidemic model

Example: 5.14 SIR model

Python: [Lecture5_5.ipynb](#)

Consider an SIR stochastic epidemic where transmissions occur at a mass-action rate of $\frac{\beta}{\kappa} * S * I$ with $\beta = 0.5$, hosts recover (gaining permanent resistance) at a rate $\gamma = 0.1$ and the total population size is $\kappa = 20$.

1. Numerically solve the system of master equations for $t = [0, 50]$ assuming that initially there are $I(0) = 5$ infections.


Let $P_{n,m}(t)$ be the probability that the system has $S(t) = n$ and $I(t) = m$ individuals present at time t . Then the process is described by the following events:

Event Name	Rate	Effect $\{\Delta n, \Delta m\}$
Transmission	$\frac{\beta}{\kappa} nm$	$\{-1, 1\}$
Recovery	γm	$\{0, -1\}$


The master equations are given by:

$$\frac{dP_{n,m}(t)}{dt} = - \left(\frac{\beta}{\kappa} nm + \gamma m \right) P_{n,m}(t) + \frac{\beta}{\kappa} (n+1)(m-1) P_{n+1,m-1} + \gamma (m+1) P_{n,m+1}$$

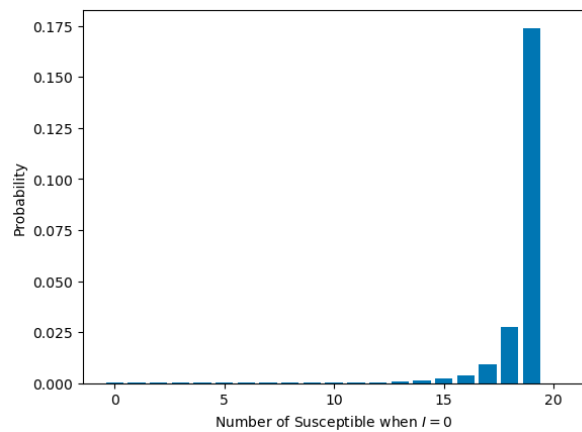
We can look at the distribution at the end of the integration period:



Or we can look at the average number of infected hosts as calculated by $\sum_{n,m} mP_{n,m}$



Or the number of susceptible hosts left when the disease goes extinct (first column of the matrix above)



Lecture 5.6 Ensemble Moment Approximations

In the previous lecture we saw how to derive a system of differential equations describing the dynamics of the state-space distribution $P_n(t)$. While providing a convenient and understandable solution in many cases, the state space $n \in \mathcal{S}$ may be very (or even infinity) large in which case approximations must be used to obtain a tractable system. Even in small or well approximated systems, subsequent interpretation of the master equation results can be cumbersome. In such cases the Ensemble Moment Approximation (EMA) can provide a useful alternative or complementary approach.

The EMA provides an approximation for the dynamics of the moments (either centered or non-centered) of the state space distribution.

Example 5.15 Moment notation

How would one express the first two raw moments and first two centered moments of the state space distribution? In a

two dimensional state-space, indicate the additional moments.

Given that the derivation of this approximation involves the use of a lot of expectations we use a shorthand bracket notation:

$$E[x] = \langle x \rangle$$

In terms of this notation the first two raw and centered moments are:

$$\begin{aligned} \mu_n(t) = \langle n \rangle &= \sum_n n P_n(t) & \text{and} & & \langle n^2 \rangle &= \sum_n n^2 P_n(t) \\ 0 & & \text{and} & & \text{Var}_n(t) &= \langle n^2 \rangle - \mu_n(t)^2 \end{aligned}$$

If the state space was 2 dimensional (with dimensions n and m) there would be four raw moments

$$\begin{aligned} \mu_n(t) = \langle n \rangle &= \sum_{n,m} n P_{n,m}(t) & \text{and} & & \langle n^2 \rangle &= \sum_{n,m} n^2 P_{n,m}(t) \\ \mu_m(t) = \langle m \rangle &= \sum_{n,m} m P_{n,m}(t) & \text{and} & & \langle m^2 \rangle &= \sum_{n,m} m^2 P_{n,m}(t) \\ \text{Var}_n(t) &= \langle n^2 \rangle - \mu_n(t)^2 & \text{and} & & \text{Var}_m(t) &= \langle m^2 \rangle - \mu_m(t)^2 \\ \text{Cov}_{n,m}(t) &= \langle nm \rangle - \mu_n(t)\mu_m(t) \end{aligned}$$

To derive the EMA we use the simple relationship that

$$\frac{d}{dt} \langle x \rangle = \left\langle \sum_{e \in \text{events}} \lambda_e \Delta_e x \right\rangle$$

where the sum is over all the possible events that can occur, λ_e is the rate at which event e occurs and $\Delta_e x$ is the effect of event e on the quantity x . Here x may be an element of the state space $x = \{n, m\}$ or a product of elements of the state space (e.g., $x = \{n^2, nm, m^2\}$).

Example 5.16 Logistic Population Growth

The logistic model is an ecological model of population growth capturing intraspecific competition. In the deterministic population limit it is given by:

$$\frac{dN}{dt} = rN \left(1 - \frac{N}{K} \right)$$

where r is the intrinsic growth rate and K is the carrying capacity.

1. Specify a corresponding CTDS-MC of logistic population growth.

While the above equation is perfectly legitimate ODE model, there is no such biological process as "intrinsic growth" or "carrying capacity", so we need to rewrite the above equation in a way that reflects its emergence from biological processes: namely births and deaths.

Important: there is no one way of doing this properly. In fact there are at least three different models of births and deaths that will ultimately result in the above ODE. But for now let's just pick one of these options.

Let's assume that births occur at a constant per-capita rate bN but deaths occur at a density-dependent rate $d(1 + \gamma N)N$ where the value of α describes the rate at which the death rate increases due to additional competition. Note that if $N \rightarrow 0$ then the death rate becomes dN .

Given this model of birth and death we have:

$$\frac{dN}{dt} = bN - d(1 + \gamma N)N$$

Discussion: What is the biological interpretation of γ ?

To equate the b , d , and γ in this new model to the original logistic growth equation we expand both equations and equate like terms (powers of N).

$$\underbrace{rN - \frac{rN^2}{K}}_{\text{logistic}} = \underbrace{(b-d)N - d\gamma N^2}_{\text{DD birth-death}}$$

So $r = (b - d)$ and $K = \frac{b-d}{d\gamma}$

Discussion: Does this equivalency make sense? How does K change and γ change?

With the new model we can specify the dynamics of our single state variable (population size $N \in \mathbb{N}$) with a 2-event CTDSMC

Event	Rate: λ_e	$\Delta_e N$
Birth	bN	+1
Death	$dN(1 + \gamma N)$	-1

2. Use the EMA to derive the dynamics of the expected population size through time.

$$\begin{aligned} \frac{d}{dt} \langle N \rangle &= \left\langle \sum_{e \in \text{events}} \lambda_e \Delta_e N \right\rangle \\ &= \langle bN * 1 + dN(1 + \gamma N) * (-1) \rangle \end{aligned}$$

Recall from expectation rules that when we take an expectation the easiest way to start simplifying is to start expanding the terms.

$$\frac{d}{dt} \langle N \rangle = \langle (b-d)N - d\gamma N^2 \rangle$$

The b , d and γ are constants so we can factor them out and split the expectation on the sum/difference to get:

$$\frac{d}{dt} \langle N \rangle = (b-d) \langle N \rangle - d\gamma \langle N^2 \rangle$$

Note that this equation has a VERY similar form as our original ODE except now we have $\langle N \rangle$ instead of N .

Also note that our ODE for the first moment $\langle N \rangle$ depends on both the first moment ($(b-d) \langle N \rangle$ term) and the second moment ($-d\gamma \langle N^2 \rangle$ term). Which means that in order to derive the dynamics for the mean we need an ODE describing $\langle N^2 \rangle$

3. Derive the ODE for $\langle N^2 \rangle$.

We use the same recipe:

$$\frac{d}{dt} \langle N^2 \rangle = \left\langle \sum_{e \in \text{events}} \lambda_e \underbrace{\Delta_e N^2}_{N_f^2 - N_0^2 \neq (N_f - N_0)^2} \right\rangle$$

where $\Delta_e N^2$ is the change in the square of the state variable NOT the change in the state variable squared.

$$\Delta_e N^2 = N_f^2 - N_0^2 \neq (N_f - N_0)^2 = (\Delta_e N)^2$$

So let's add to our table

Event	Rate: λ_e	$\Delta_e N$	$\Delta_e N^2$
Birth	bN	$N + 1 - N = 1$	$(N + 1)^2 - N^2 = 2N + 2$
Death	$dN(1 + \gamma N)$	$N - 1 - N = -1$	$(N - 1)^2 - N^2 = -2N + 2$

Event	Rate: λ_e	$\Delta_e N$	$\Delta_e N^2$
Substituting in these values into our ODE we have:			

$$\begin{aligned} \frac{d}{dt} \langle N^2 \rangle &= \langle bN(2N+2) + dN(1+\gamma N)(-2N+2) \rangle \\ &= \langle 2N^2(b+(\gamma-1)d) + 2N(b+d) - 2\gamma dN^3 \rangle \\ &= 2(b+(\gamma-1)d) \langle N^2 \rangle + 2(b+d) \langle N \rangle - 2\gamma d \langle N^3 \rangle \end{aligned}$$

So this equation depends on the third raw moments. This is a problem!!! What do we do?

4. Now consider the Lotka-Volterra model of inter-specific competition between two species. Derive the corresponding EMA for the mean dynamics of this model.

$$\begin{aligned} \frac{dN_1}{dt} &= b_1 N_1 - d_1 N_1 (1 + \gamma_{1,1} N_1 + \gamma_{1,2} N_2) \\ \frac{dN_2}{dt} &= b_2 N_2 - d_2 N_2 (1 + \gamma_{2,1} N_1 + \gamma_{2,2} N_2) \end{aligned}$$

Take Home Challenge: Derive the relationship between b_i , d_i and γ_i and the parameters of the classic Lotka-Volterra model:

$$\begin{aligned} \frac{dN_1}{dt} &= r_1 N_1 \left(1 - \frac{N_1 + \alpha_2 N_2}{K_1} \right) \\ \frac{dN_2}{dt} &= r_2 N_2 \left(1 - \frac{N_2 + \alpha_1 N_1}{K_2} \right) \end{aligned}$$

Here we have a 4 event CTDS process with two state variables N_1 and N_2

Event	Rate: λ_e	$\Delta_e N$
Birth of 1	$b_1 N_1$	$\{+1, 0\}$
Death of 1	$d_1 N_1 (1 + \gamma_{1,1} N_1 + \gamma_{1,2} N_2)$	$\{-1, 0\}$
Birth of 2	$b_2 N_2$	$\{0, +1\}$
Death of 2	$d_2 N_2 (1 + \gamma_{2,1} N_1 + \gamma_{2,2} N_2)$	$\{0, -1\}$

So

$$\begin{aligned} \frac{d \langle N_1 \rangle}{dt} &= \langle b_1 N_1 * 1 + d_1 N_1 (1 + \gamma_{1,1} N_1 + \gamma_{1,2} N_2) * (-1) \rangle \\ \frac{d \langle N_2 \rangle}{dt} &= \langle b_2 N_2 * 1 + d_2 N_2 (1 + \gamma_{2,1} N_1 + \gamma_{2,2} N_2) * (-1) \rangle \end{aligned}$$

Take Home Challenge: Simplify the ODEs above using expectation rules. What higher order moments do these first moments depend on?

Moment Closure

As exemplified above, when there is non-linearity in the system the dynamics of the first moments will depend on the second raw moments. So in order to solve the problem with a small number of ODEs we need to make an assumption.

Rule 1: Assume that the n^{th} central moment is of order $\mathcal{O}(\epsilon^{n-1})$.

So in this case $Var(x) \approx \epsilon$, $Skew(x) \approx \epsilon^2$ etc.

Rule 2: If we want to understand the dynamics of the n^{th} central moment then we need to approximate to order $\mathcal{O}(\epsilon^n)$ assuming the $(n + 2)^{\text{nd}}$ moment is 0.

So if we want to understand the dynamics of the mean μ_x , we have to model the mean and the variance and we assume the skew is 0. If we want to understand the mean and the variance we have to model the mean, variance, and skew and assume the kurtosis is 0. So you always include one more moment than you want to approximate well.

From a practical standpoint, if we want to understand the dynamics of the mean we are effectively modelling the dynamics of the population with a gaussian (a distribution that has a mean and variance but no skew).

Important Note: Just because the central moment is small $\langle (X - \mu)^n \rangle \approx 0$ DOES NOT imply the raw moment is small $\langle X^n \rangle \approx 0$

Example 5.17 Logistic Population Growth cont

Python: [Lecture5_6.ipynb](#)

Suppose we want to understand how stochasticity impacts mean population size.

1. Derive an EMA for the mean dynamics (including the second moment and closing at the third moment)

So we want to assume that $\text{Skew}(N) \approx 0$

From above we have:

$$\begin{aligned} \frac{d}{dt} \langle N \rangle &= (b - d) \langle N \rangle - d\gamma \langle N^2 \rangle \\ \frac{d}{dt} \langle N^2 \rangle &= 2(b + (\gamma - 1)d) \langle N^2 \rangle + 2(b + d) \langle N \rangle - 2\gamma d \langle N^3 \rangle \end{aligned}$$

Take Home Challenge: Using the moment rules from Topic 1 show that:

$$\langle x^3 \rangle = \text{Skew}(x) + 3 \langle x^2 \rangle \langle x \rangle - 2 \langle x \rangle^3$$

or given our approximation that $\text{Skew}(x) \approx 0$

$$\langle x^3 \rangle = 3 \langle x^2 \rangle \langle x \rangle - 2 \langle x \rangle^3$$

$$\begin{aligned} \frac{d}{dt} \langle N \rangle &= (b - d) \langle N \rangle - d\gamma \langle N^2 \rangle \\ \frac{d}{dt} \langle N^2 \rangle &= 2(b + (\gamma - 1)d) \langle N^2 \rangle + 2(b + d) \langle N \rangle - 2\gamma d \left(3 \langle N^2 \rangle \langle N \rangle - 2 \langle N \rangle^3 \right) \end{aligned}$$

We now have 2 ODEs for 2 variables $\langle N \rangle$ and $\langle N^2 \rangle$

2. Evaluate the EMA above numerically. Plot the deterministic expectation, the mean dynamics, and plus or minus 1 standard deviation of this mean. Under what conditions is the approximation valid?

We have to choose parameters. Let's assume that $b = 0.1$ and $d = 0.05$ and $\gamma = 0.005$. In this way the carrying capacity of the system is: $\frac{b-d}{\gamma d} = 200$

If we start with exactly 5 individuals in the population then $\langle N \rangle(0) = 20$ and $\text{Var}(N(0)) = \langle N^2 \rangle(0) - (\langle N \rangle(0))^2 = 0$ (there is no initial variability).

Hence:

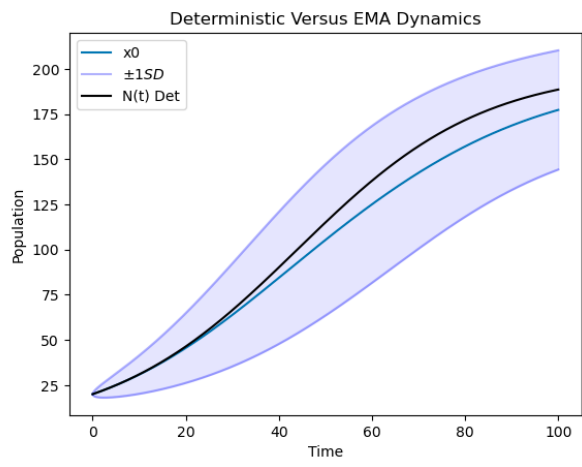
$$\langle N^2 \rangle(0) = 20^2$$

To obtain the variance from $\langle N \rangle$ and $\langle N^2 \rangle$ we have:

$$\text{Var}(N) = \langle N^2 \rangle - \langle N \rangle^2$$

The results are roughly valid for this parameter condition. Note that if we start with far fewer individuals $N(0) = 2$ then a lot of the trajectories goes extinct. This dramatically increases the variance and violates our assumption.

3. Compare the results of the EMA to the deterministic solution.




Topic 5c: Continuous Time/Discrete State

Lecture 5.7 Diffusion Approximations

Review: WF Model

Recall from topic 2 that the WF model is a DTDSMC describing neutral genetic drift in a small population. We can depict the dynamics of the WF model using the following schematic:



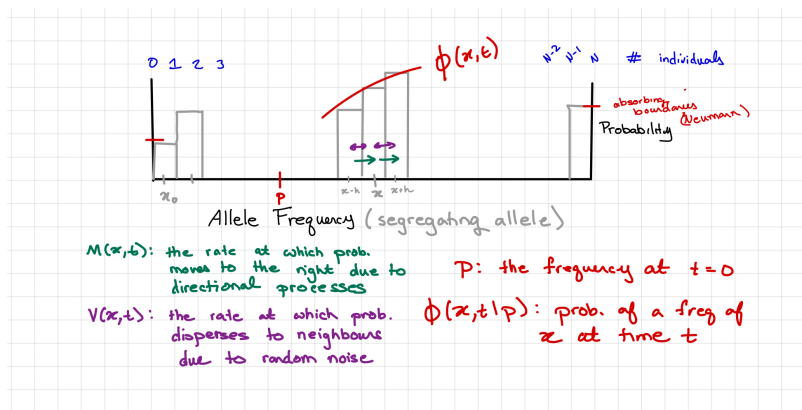
The probability of transitioning between having i 'A' parents to j 'A' offspring is:

$$P(X_{t+1} = j | X_t = i) = p_{ij} = \binom{N}{j} \frac{i^j}{N^j} \left(\frac{N-i}{N}\right)^{N-j} = \underbrace{\binom{N}{j} p^j (1-p)^{N-j}}_{B_j(N,p)}$$

When the total population size N is reasonably small we can easily simulate this process, evaluate the dynamics numerically, and solve for the time to absorption/probability of absorption. But as N gets bigger these things become harder to do as the transition probability matrix gets too big.

The diffusion approximation is an approximation of the WF model for large (but still finite!) populations. The result of the diffusion approximation is a PDE of the function $\phi(x, t | p_0)$ describing the probability that a population has a (segregating) allele frequency x at time t given that the population started with an allele frequency of p_0 . This ODE is so widely used that it has its own name the "Kolmogorov Forward Equation".

Deriving the Kolmogorov Forward Equation



To derive the diffusion approximation to the Wright Fisher model we begin by assuming that the number of individuals is large such that the probability that there are $0 < i < N$ individuals in the population with the 'A' allele can be approximated as a continuous frequency $0 < x < 1$. Note that space is subdivided into rectangles of width $\frac{1}{N}$ which are centred offset from the counts of individuals. We do not then consider the extreme cases of $i = 0$ and a $i = N$ and hence this approximation gives us the distribution of segregating allele frequencies.

Let $\phi(x, t | p)$ then be the probability that a 'A' allele that is segregating in the population at time has frequency x at time t given that the frequency starts at p at time 0. Probability can move either due to directional processes (e.g., selection, migration,

mutation) or **diffuse** in an arbitrary direction due to random drift.

Let's define:

- $m(x, t)$ as the probability that the process moves to the right due to directional processes
- $v(x, t)$ the probability that the process moves to the left ($\frac{1}{2}$ of the time) or to the right ($\frac{1}{2}$ of the time) due to random processes.

We can write the change in the volume of probability as:

$$\begin{aligned} \frac{\phi(x, t + \Delta t|p)\Delta x - \phi(x, t|p)\Delta x}{\Delta t} &= \phi(x, t|p)\Delta x - \phi(x, t|p)\Delta x \\ &- \left(m(x, t)\frac{\Delta t}{\Delta t} + v(x, t)\frac{\Delta t}{\Delta t} \right) \phi(x, t|p)\Delta x \\ &+ \frac{1}{2}v(x - \Delta x)\frac{\Delta t}{\Delta t}\phi(x - \Delta x, t|p)\Delta x \\ &+ \frac{1}{2}v(x + \Delta x)\frac{\Delta t}{\Delta t}\phi(x + \Delta x, t|p)\Delta x \\ &+ m(x - \Delta x)\frac{\Delta t}{\Delta t}\phi(x - \Delta x, t|p)\Delta x \end{aligned}$$

Let

- $M(x, t) = E[\Delta x]$ be the expected amount of movement to the right

$$M(x, t) = m(x, t) * \Delta x$$

- $V(x, t) = \text{Var}[\Delta x] = E[\Delta x^2] - E[\Delta x]^2$ be the variance in the movement due to random processes.

$$\begin{aligned} V(x, t) &= \left(\frac{v}{2}\Delta x^2 + \frac{v}{2}(-\Delta x)^2 \right) - \underbrace{\left(\frac{v}{2}\Delta x + \frac{v}{2}(-\Delta x) \right)^2}_{=0} = \Delta x^2 v(x, t) \\ \frac{d\phi(x, t)}{dt} \frac{\Delta x}{\Delta x} &= - \left(\frac{\frac{M(x, t)}{\Delta x}\phi(x, t|p)\Delta x - \frac{M(x - \Delta x, t)}{\Delta x}\phi(x - \Delta x, t|p)\Delta x}{\Delta x} \right) \\ &+ \frac{1}{2} \frac{\frac{V(x + \Delta x, t)}{\Delta x^2}\phi(x + \Delta x, t|p)\Delta x - \frac{V(x, t)}{\Delta x^2}\phi(x, t|p)\Delta x}{\Delta x} \\ &+ \frac{1}{2} \frac{\frac{V(x - \Delta x, t)}{\Delta x^2}\phi(x + \Delta x, t|p)\Delta x - \frac{V(x, t)}{\Delta x^2}\phi(x, t|p)\Delta x}{\Delta x} \end{aligned}$$

The first row is the can be simplified using the definition of the one-step derivative. The second two rows can be simplified using the definition of the second derivative

$$\frac{d\phi(x, t)}{dt} = - \frac{\partial M(x, t)\phi(x, t|p)}{\partial x} + \frac{1}{2} \frac{\partial^2 V(x, t)\phi(x, t|p)}{\partial x^2}$$

Where the first term captures the change in the allele frequency due to directional processes (this is known as *mathematical drift*) and $M(x, t)$ is the infinitesimal mean.

The second term captures the change in the allele frequency due to random processes (this is known as *mathematical diffusion* or *biological drift*) and $V(x, t)$ is the **infinitesimal variance**.

The Wright-Fisher Model in Review

- The Wright-Fisher Model is a DTDS model of neutral genetic drift capturing the change in allele frequencies from generation to generation due to the random sampling (of parents).
- The WF model can be used to follow the number/frequency of a 'A' allele in a finite population.
- In the absence of mutation the WF model has two absorbing states: fixation of the 'A' and loss of the 'a'. We can use first-step analysis to calculate the relative probabilities of ending of in each of these absorbing states.

- In the presence of mutation, we can derive the stationary distribution of the WF model capturing the amount of time the population stays in a particular state in the long term.
- The Moran model is an alternative to the WF model capturing evolution in approximately continuous time. There are notable differences (by a factor of 2) in the rate of genetic drift between the WF and Moran models. Whether we model in discrete or continuous time matters!!
- The Kingman coalescent uses the WF model to describe the genealogical ancestry of a sample of genes from the present day.
- The Kolmogorov forward equation is a PDE describing the dynamics of the WF model in a large but finite population.
-

Lecture 5.8 Birth-Death Models

- Results from the appendix of Raup 1985

The Yule Model

The Yule process is a stochastic process proposed by George Yule in 1924 to describe the distribution of species among different taxonomic groups, known as clades.

In the context of ecology or biology, the Yule process is a branching process that models the evolution of species in a phylogenetic tree. The process assumes that new species can arise by speciation events, where an existing species splits into two new ones.

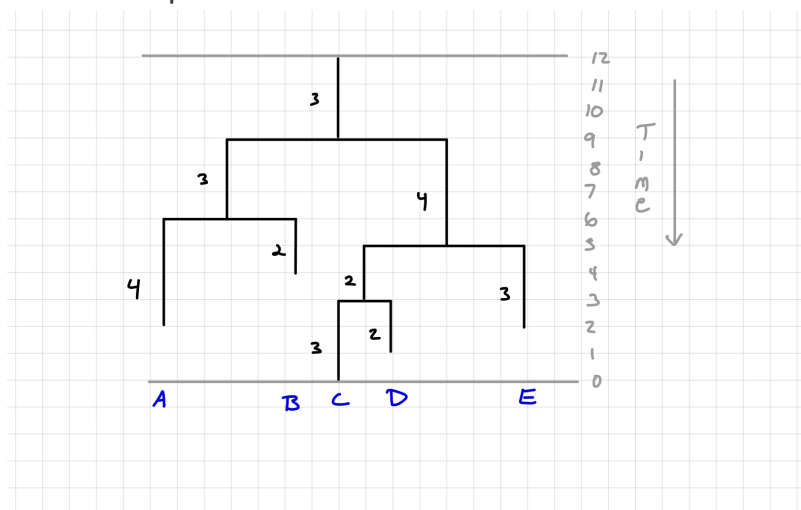
Let $N(t)$ be the number of species present in a clade at time t . Speciation occurs (species are 'born') at a constant per-capita rate λ . Then $N(t)$ is a continuous-time stochastic process known as the Yule process or "birth process".

Simulating a Phylogenetic Tree under the Yule Model

The key to the simulation here is the storage of data. We can represent a tree as a **cophenetic** matrix where element $C_{i,j}$ give the total amount of "shared ancestry between the lineages".

Example: 5.XX cophenetic matrices

What is the cophenetic matrix for the tree shown below?



To simulate a tree under the Yule process we simply need to simulate the corresponding cophenetic matrix.

Initialization: Initially the cophenetic matrix is:

Initially the time is $t = 0$

$$\mathbf{C}(0) = [0]$$

There is a single lineage/point that shares no ancestry with itself.

Iteration For ($t < t_{max}$):

- Calculate the total rate of speciation $\Lambda = \lambda * n$ where n is the number of lineages in \mathbf{C} .
- Draw the time until the next event $\Delta t \sim Exp(\Lambda)$. Update $t+ = \Delta t$
- Update the matrix by adding Δt to each diagonal element. This extends each existing branch by the amount Δt .
- Add a branching event
 - Choose a parent to branch at random
 - Copy and paste the parent's row and column and add t to the current diagonal. This adds the offspring which by definition shares t ancestry with itself.

Output: Return the cophenetic matrix

Example: 5.18 Yule Simulation

Python Lec_5.8

1. Simulate the Yule process with $\lambda = 1$ for $T = 1.5$ units of time. Sketch the corresponding tree.

2. Draw the "Lineage Through Time" plot for your simulation

The Yule simulations above have a behaviour similar to exponential population growth. How big should the tree be after T units of time and how should the number of lineages change through time?

Clade size in the Yule model


Specifically, suppose we start with 1 lineage at time $t = 0$ (the origin of the tree), how big is the tree at time $t = T$ later? To calculate this we consider the master equations for the probability that there are n lineages.

$$\frac{dP_n(t)}{dt} = -\lambda n P_n + \lambda(n-1)P_{n-1} \quad P_n(0) = \begin{cases} 0 & n \neq 1 \\ 1 & n = 1 \end{cases}$$

The solution to these ODEs are:

$$P(n, t) = e^{-n\lambda t} (e^{\lambda t} - 1)^{(n-1)}$$

The solutions for $n = 1$ (red) to $n = 5$ (blue) are shown below:



The Constant-Rate Birth-Death Model

The Yule model has one key assumption that deviates from our knowledge of the natural world, the fact that species can both speciate (be born) and go extinct (die). The simplest model that captures this process is the constant rate birth-death model in