

CS 364

Computational Biology

Sara Mathieson
Haverford College

Outline

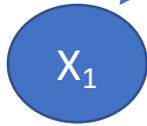
- HMM example in population genetics
- Recap Viterbi Algorithm
- Forward-Backward Algorithm
- Posterior Decoding

Notes:

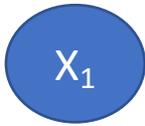
- Today and Tuesday: HMMs
- Next Thursday: review
- Please fill out exam/generative AI poll!

Hidden Markov models

Lunch on Monday



Lunch on Tuesday



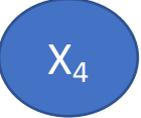
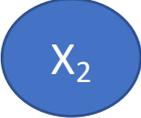
Observations



.....



Mood: "Busy" or "Relaxed"



.....



Hidden state

Observations

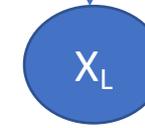
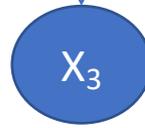
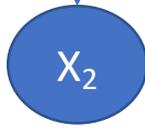
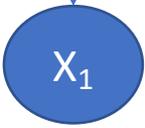
Hidden state



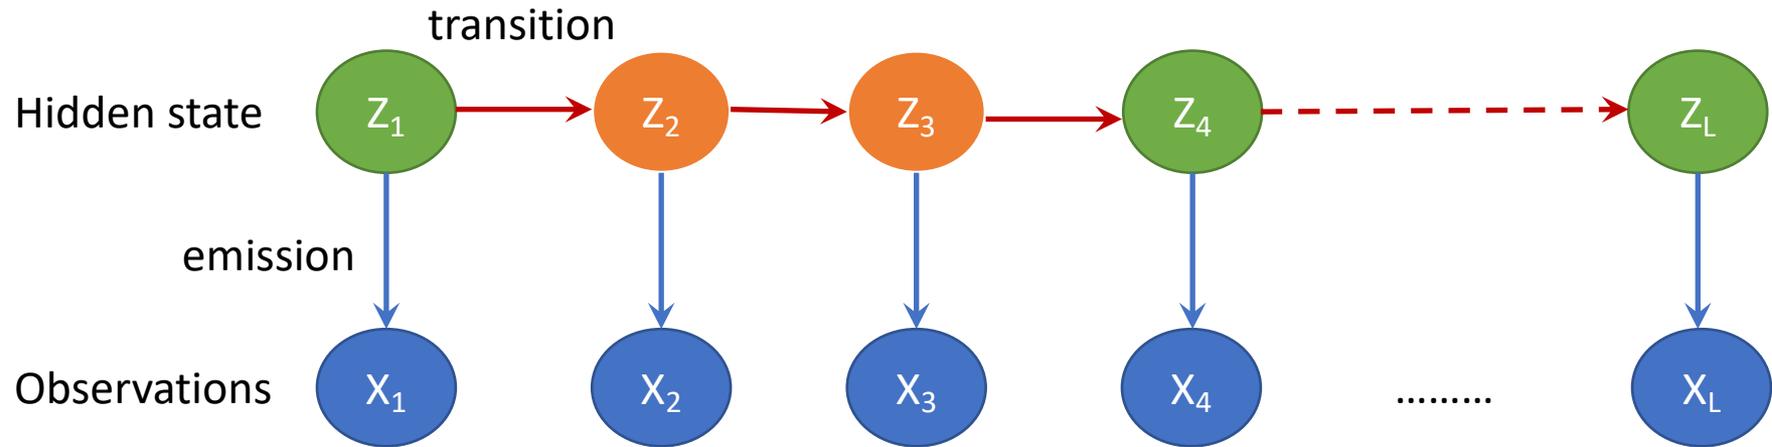
.....



Observations



Markov property: hidden state i depends on state $i-1$ but *only* on state $i-1$



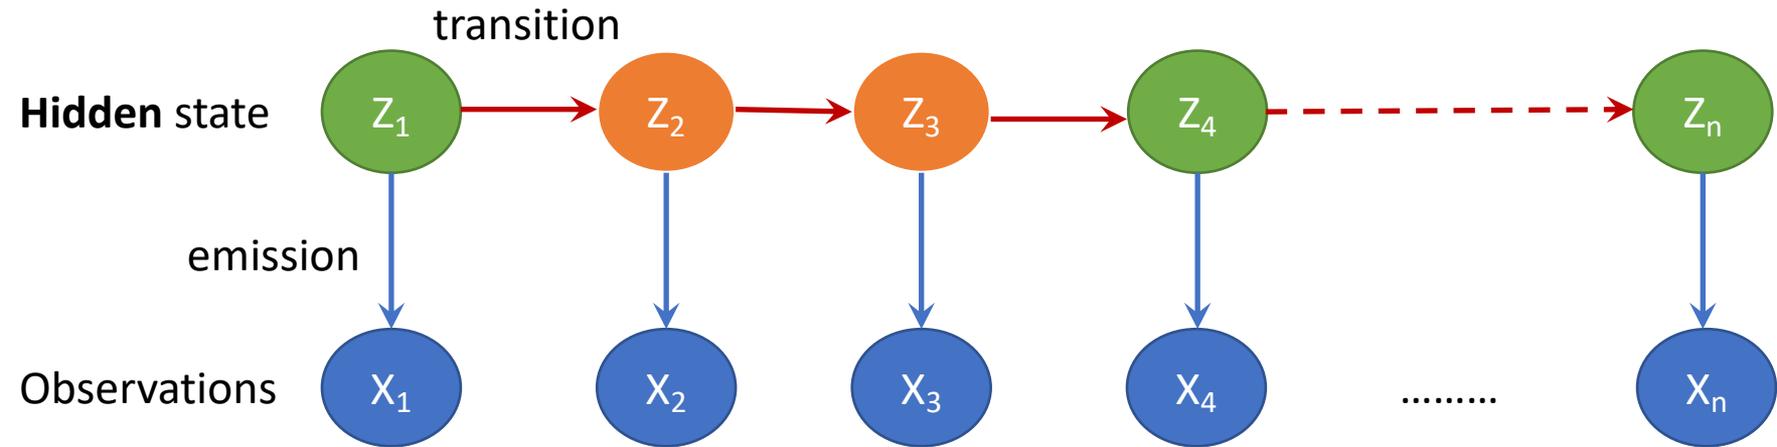
Hidden Markov model:

We have a series of **observations** that depend on some underlying **hidden state**.

Parameters: for each value of the hidden state, what's the probability of each possible observations [i.e. what is $P(X_i | Z_i)$?]. This is the **emission probability** \longrightarrow

If we are in hidden state Z_i , what is the probability that we are in each of the possible hidden states in Z_{i+1} . [i.e. what is $P(Z_{i+1} | Z_i)$?]. This is the **transition probability** \longrightarrow

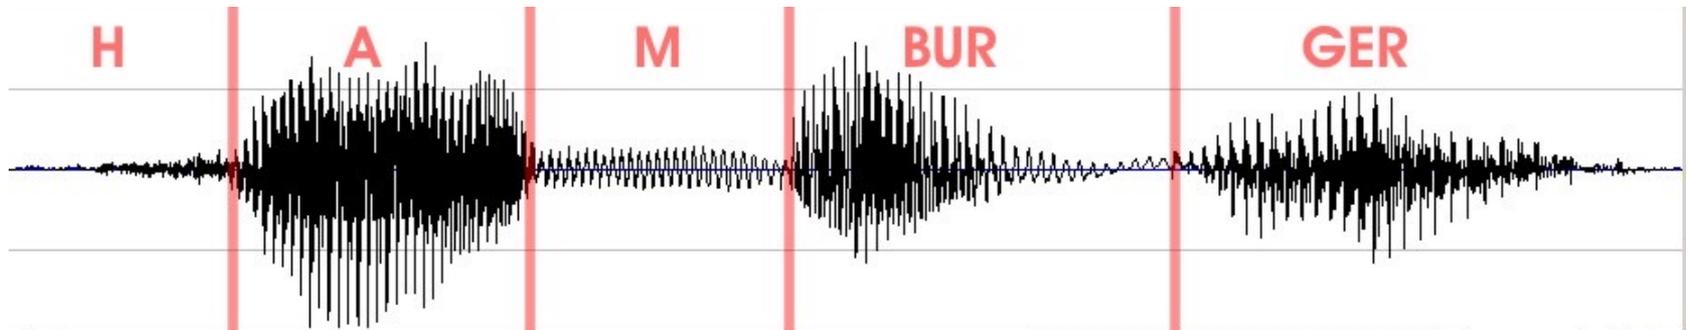
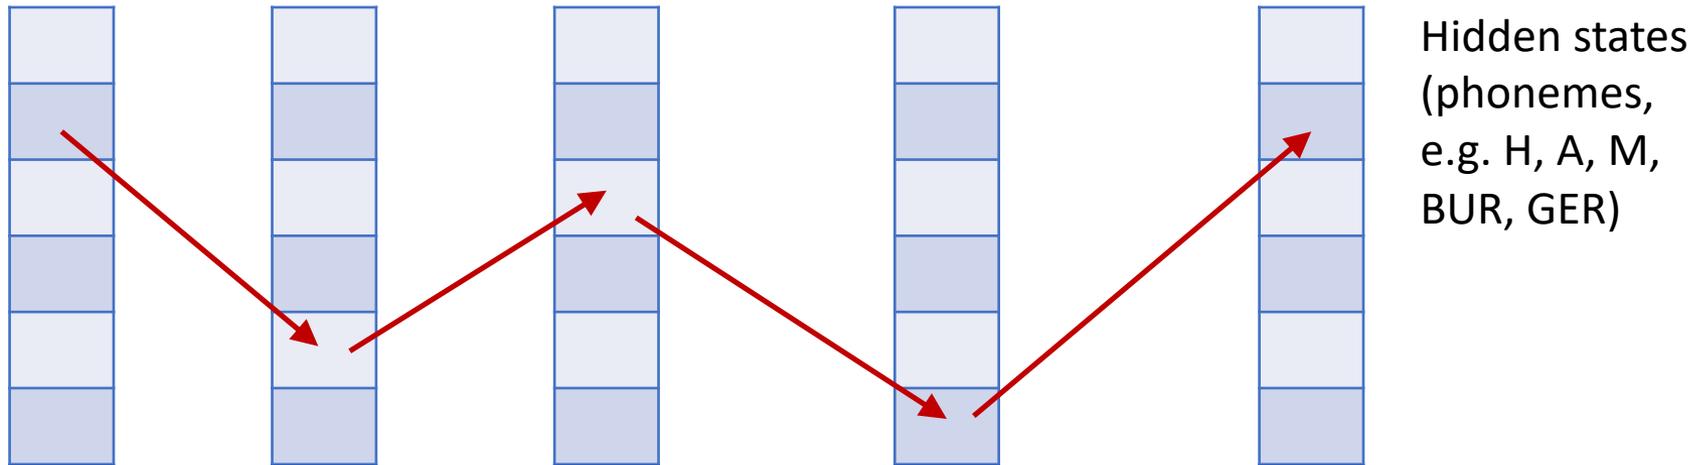
Markov property: hidden state i depends on state $i-1$ but *only* on state $i-1$



What can we learn

- 1) Given the observations, what were the hidden states? i.e. given the record of what I ate for lunch over a week, what was my mood on each day?
- 2) What are the parameters. e.g. if I am busy one day, what is the probability that I am relaxed the next day (transition)? If I am relaxed today, what is the probability that I eat out (emission)?

Example: Speech recognition



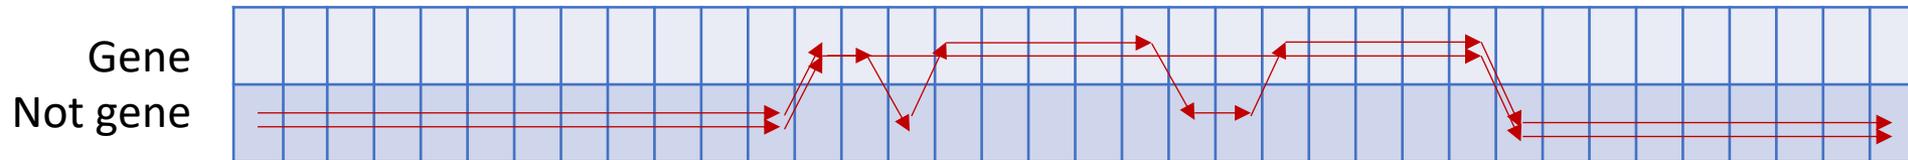
Observations (waveforms)

Example: sequences (gene finding)

Where is the gene?

Coding sequences have higher CG content than noncoding regions

ACATCAGCTACGATCGATGCGGGCATGAGTCCCATATATTTAG



Transition probabilities

	Gene	Not
Gene	0.9	0.1
Not	0.05	0.95

(e.g. average gene length is 10
and 1/20 of the genome is genic)

Emission probabilities

	A	C	G	T
Gene	0.25	0.25	0.25	0.25
Not	0.35	0.15	0.15	0.35

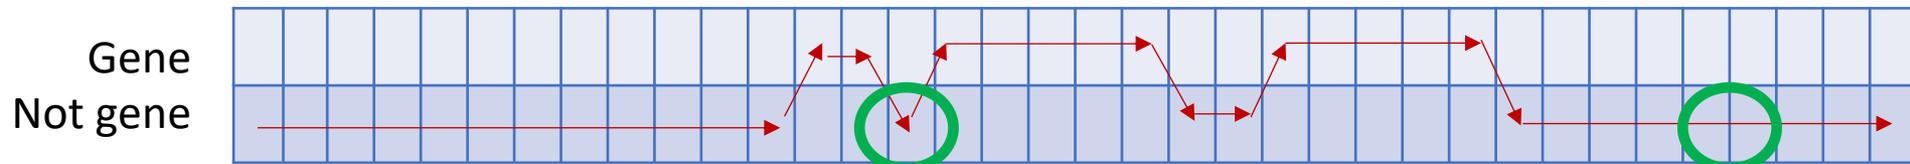
(e.g. GC content is 50% in genes and 30% outside)

Example: sequences (gene finding)

Where is the gene?

Coding sequences have higher CG content than noncoding regions

ACATCAGCTACGATCGATGCGGGGCATGAGTCCCATATATTTAG



The most likely path says that this nucleotide is not in a gene

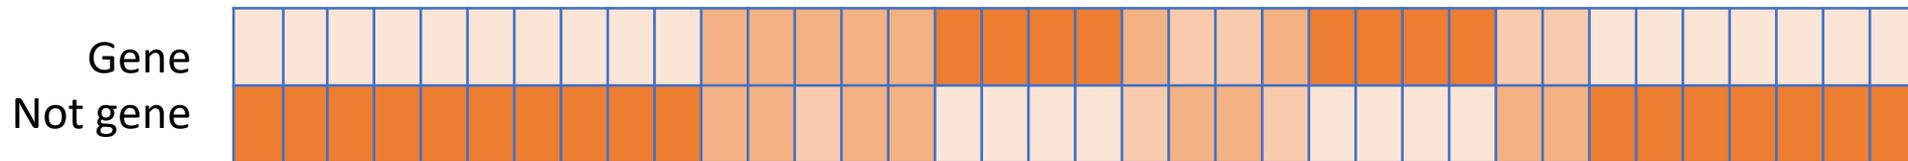
But surely it is more likely to be in a gene than this one?

Example: sequences (gene finding)

Where is the gene?

Coding sequences have higher CG content than noncoding regions

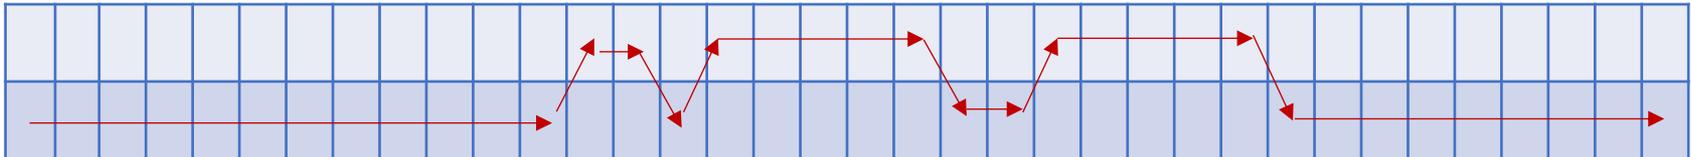
ACATCAGCTACGATCGATGCGGGCATGAGTCCCATATATTTAG



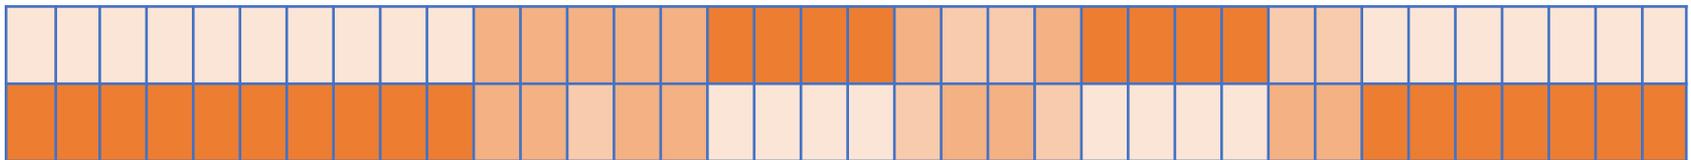
Find probability of each hidden state, rather than most likely path

Two ways of learning the hidden states

a) Find the most likely path: Viterbi algorithm



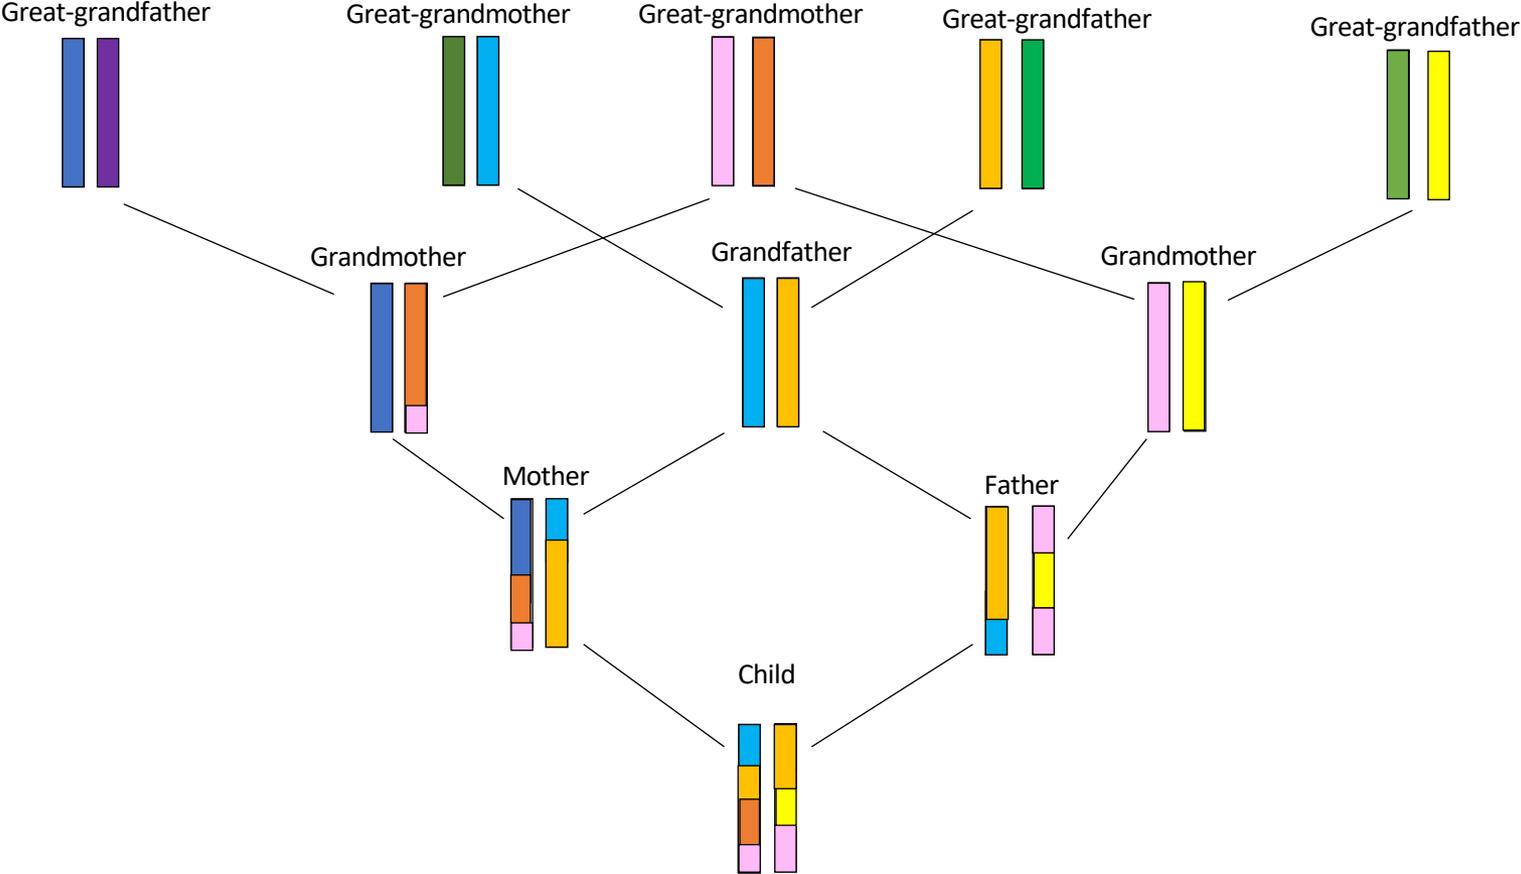
b) Find posterior probability of each state: Forward-backward algorithm



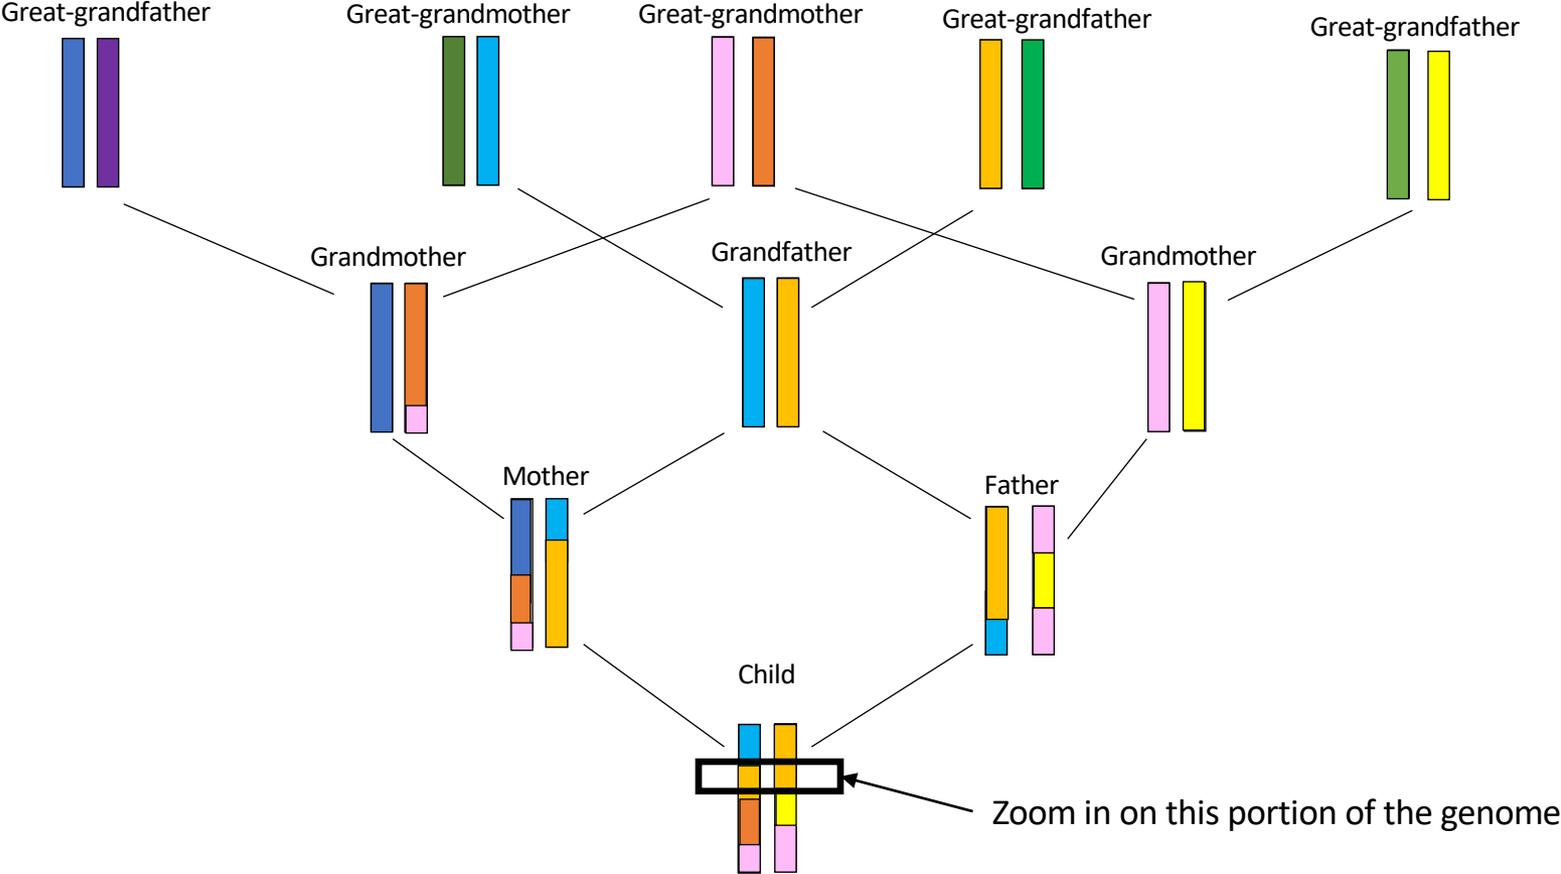
Note that the Viterbi path is NOT the same as taking the maximum cells from the Forward-Backward algorithm (“posterior decoding”). The Viterbi path must be a path that is actually possible, but the maximum posterior might not be (e.g. because it implies impossible transitions).

HMM example from
population genetics

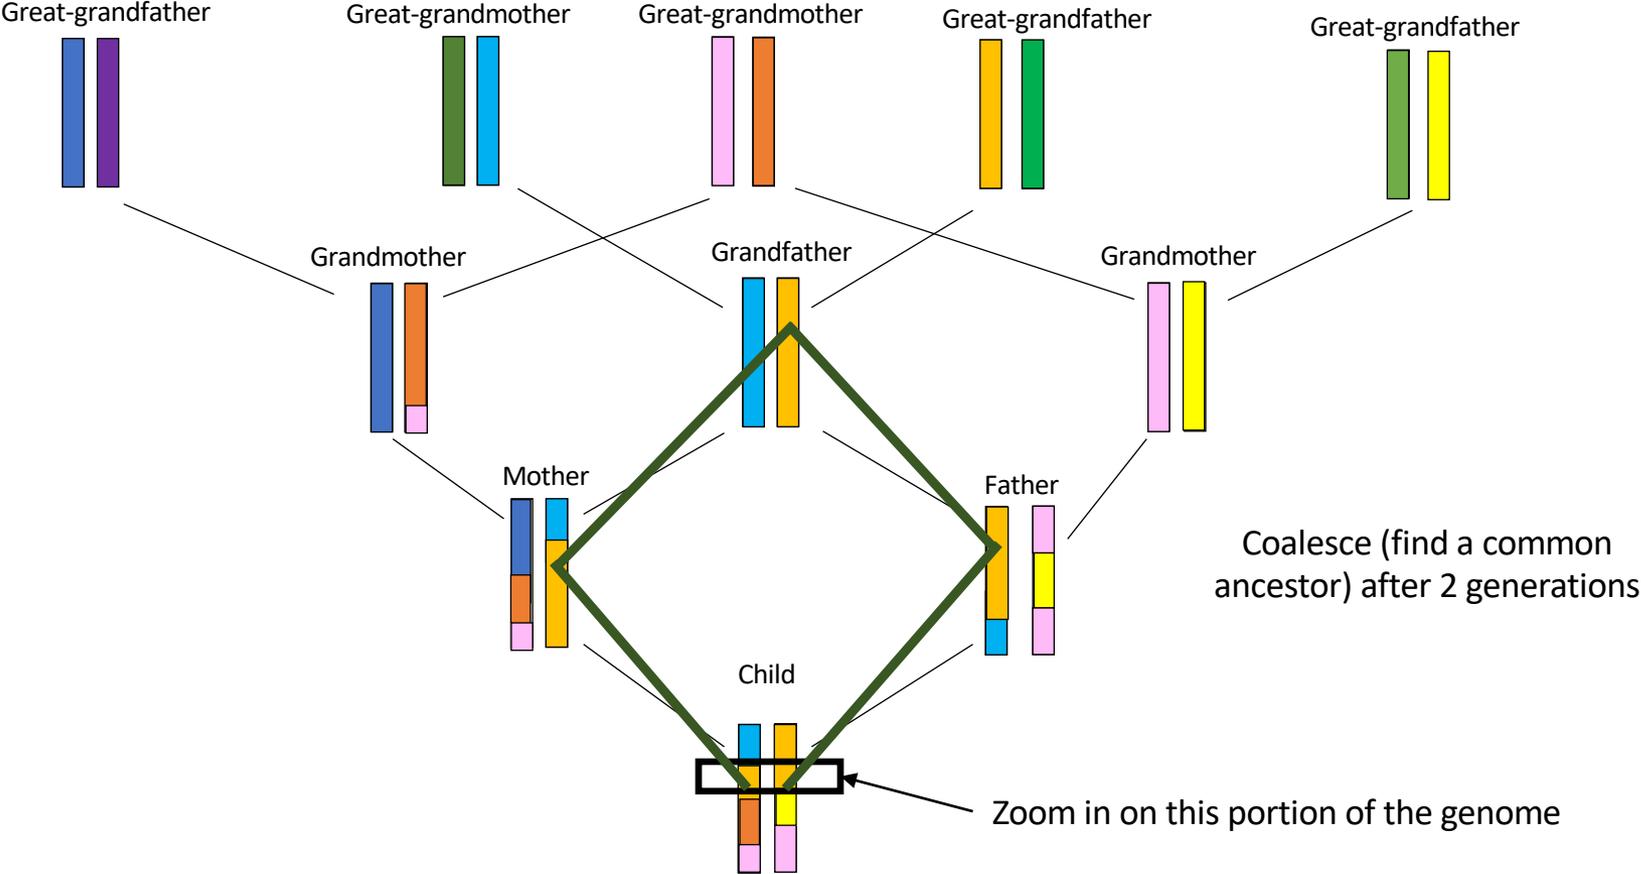
Recombination over time



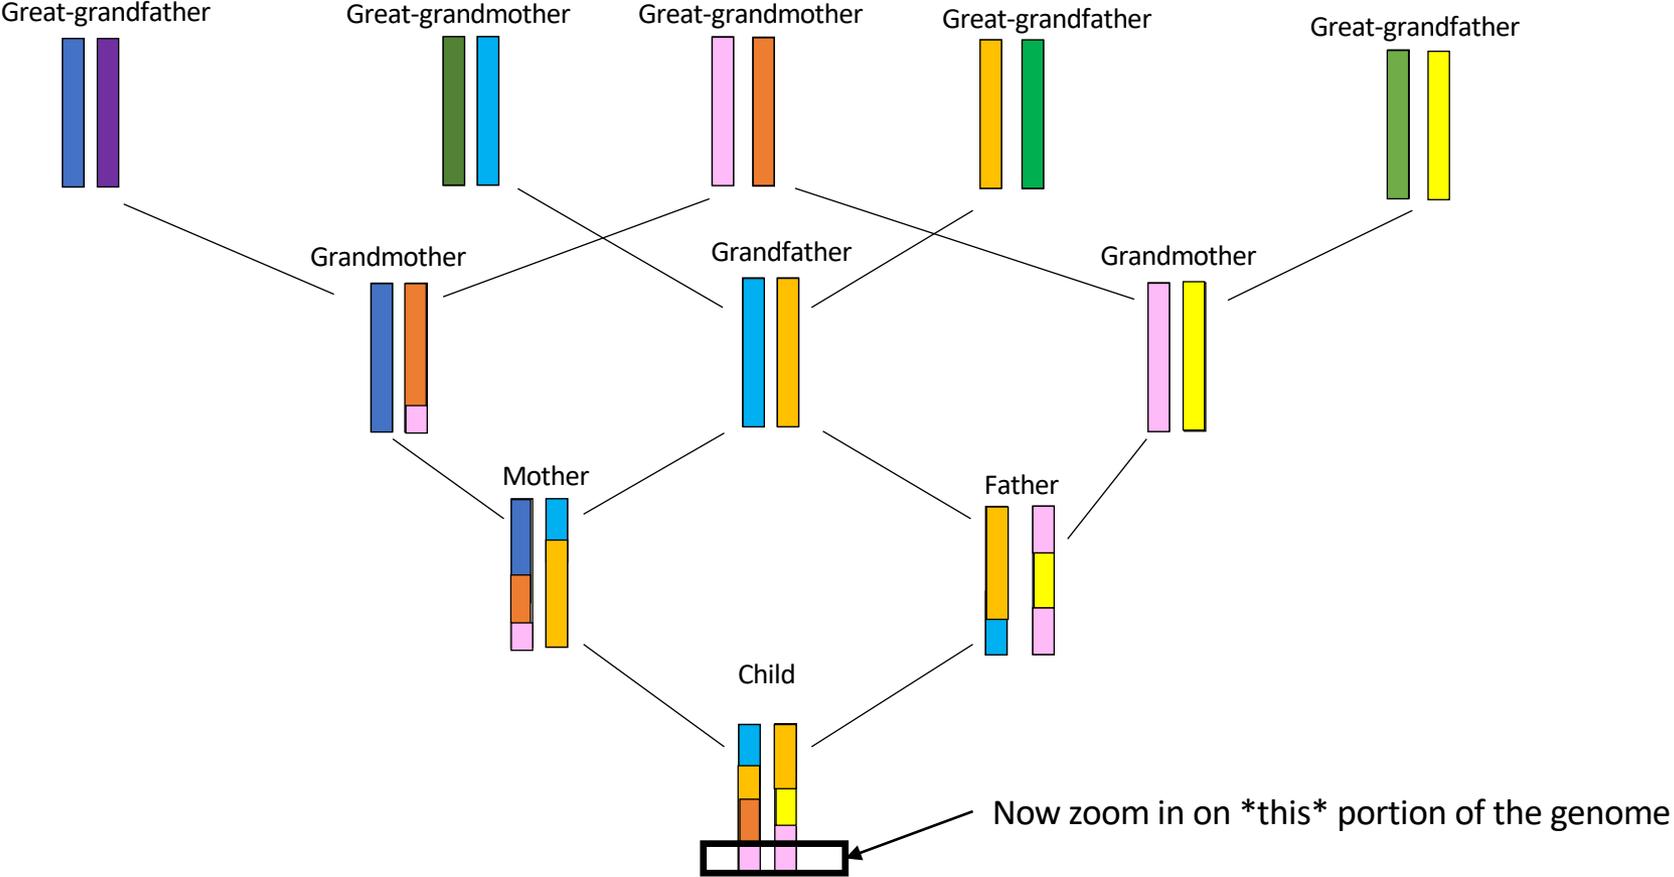
Recombination over time



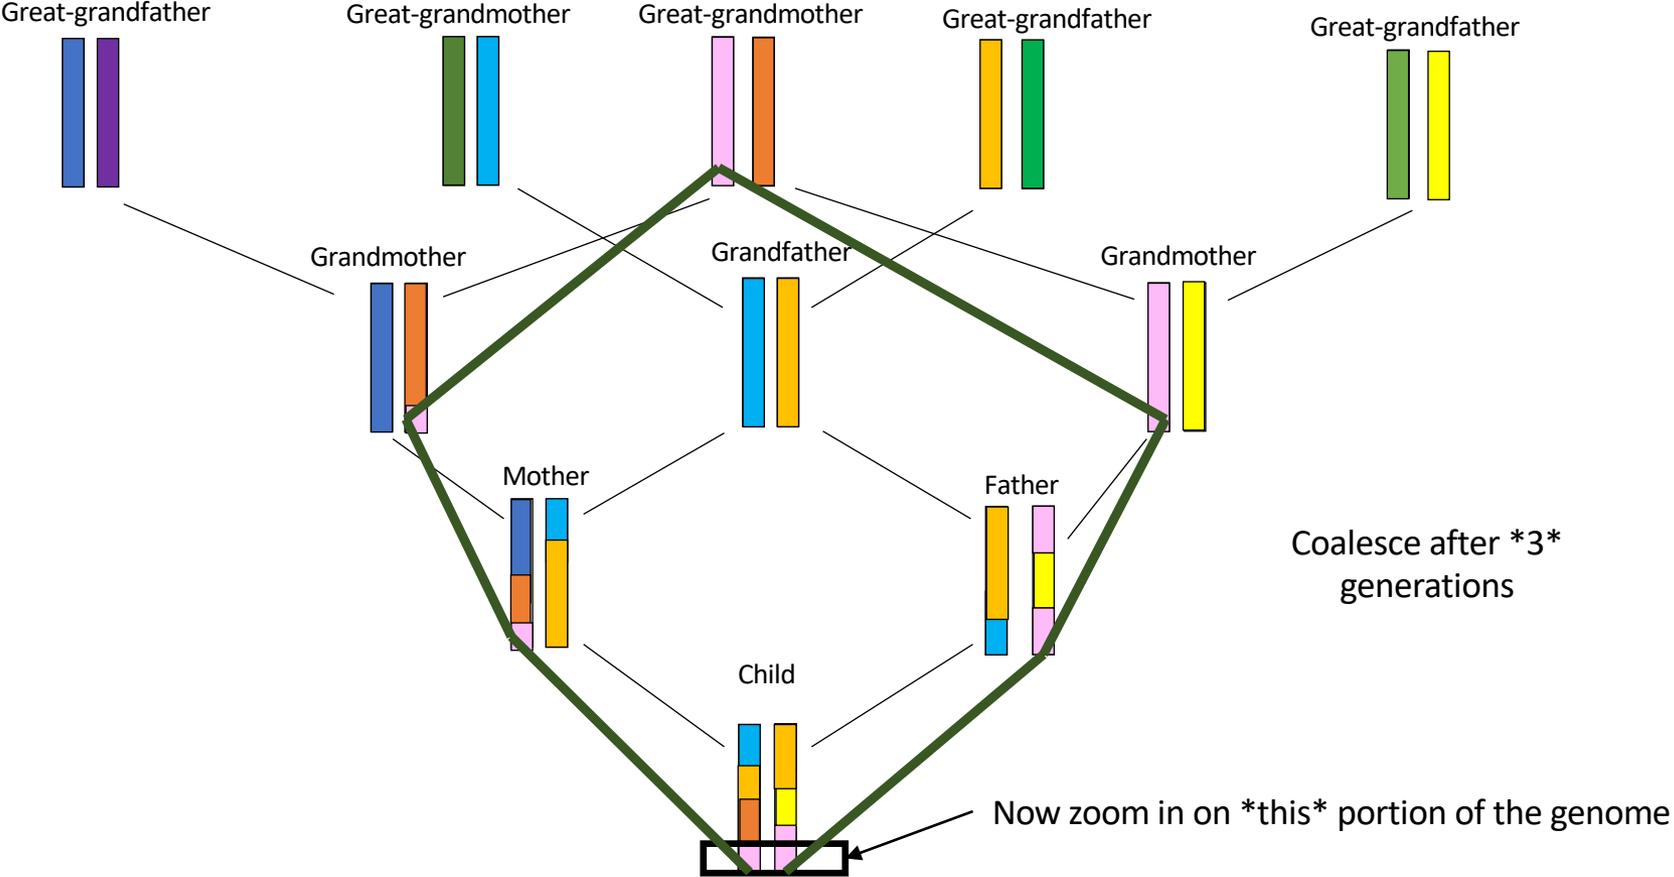
Recombination over time



Recombination over time



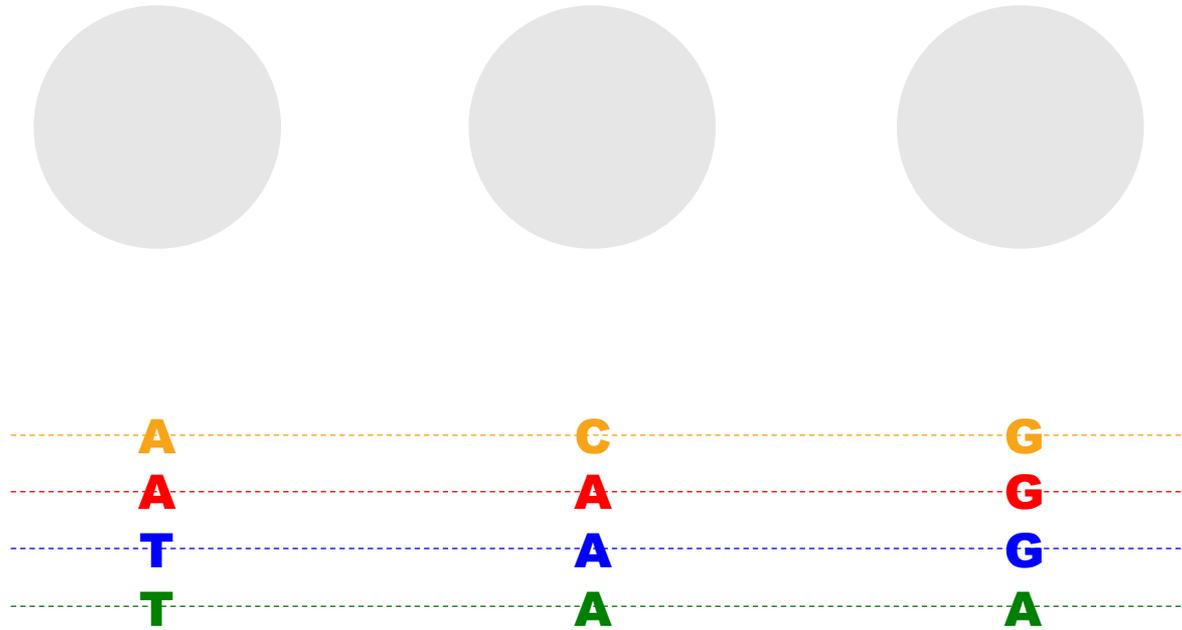
Recombination over time



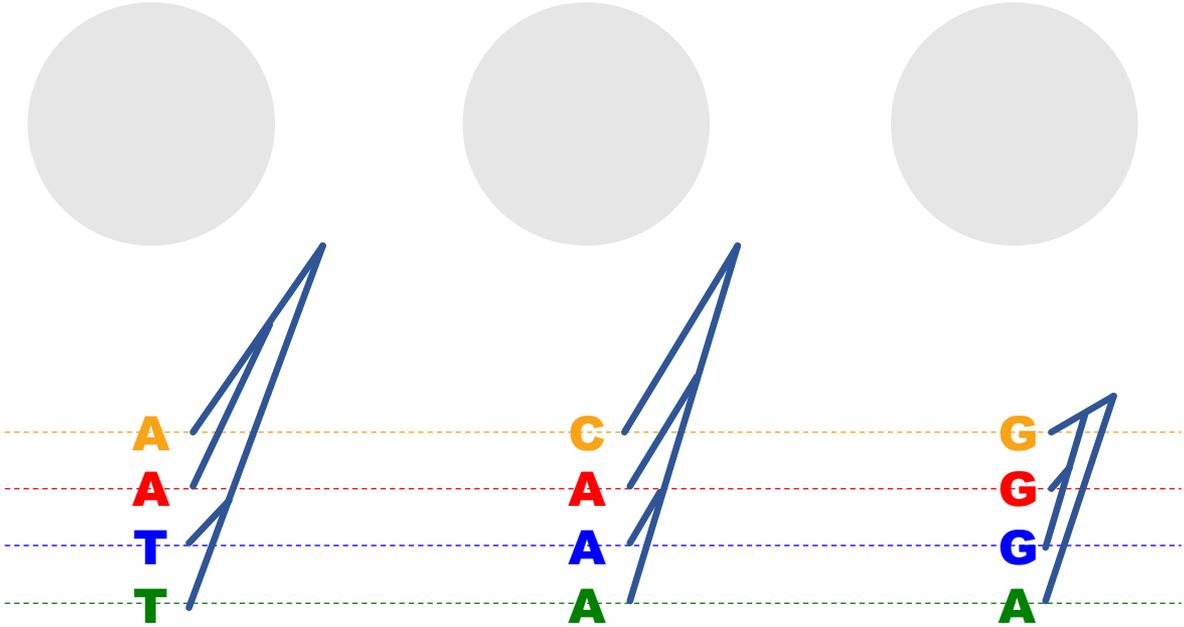
How could we encode this as an HMM?

- Take-home message: the tree changes across the genome! Both topology (for $n > 2$) and branch lengths

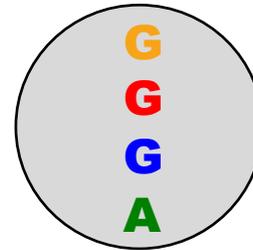
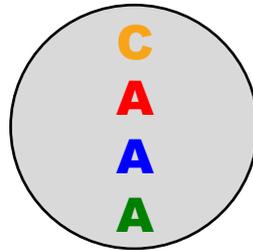
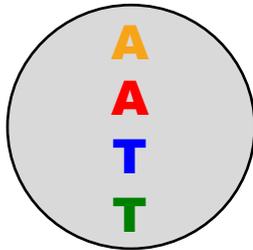
Sequence data at many sites



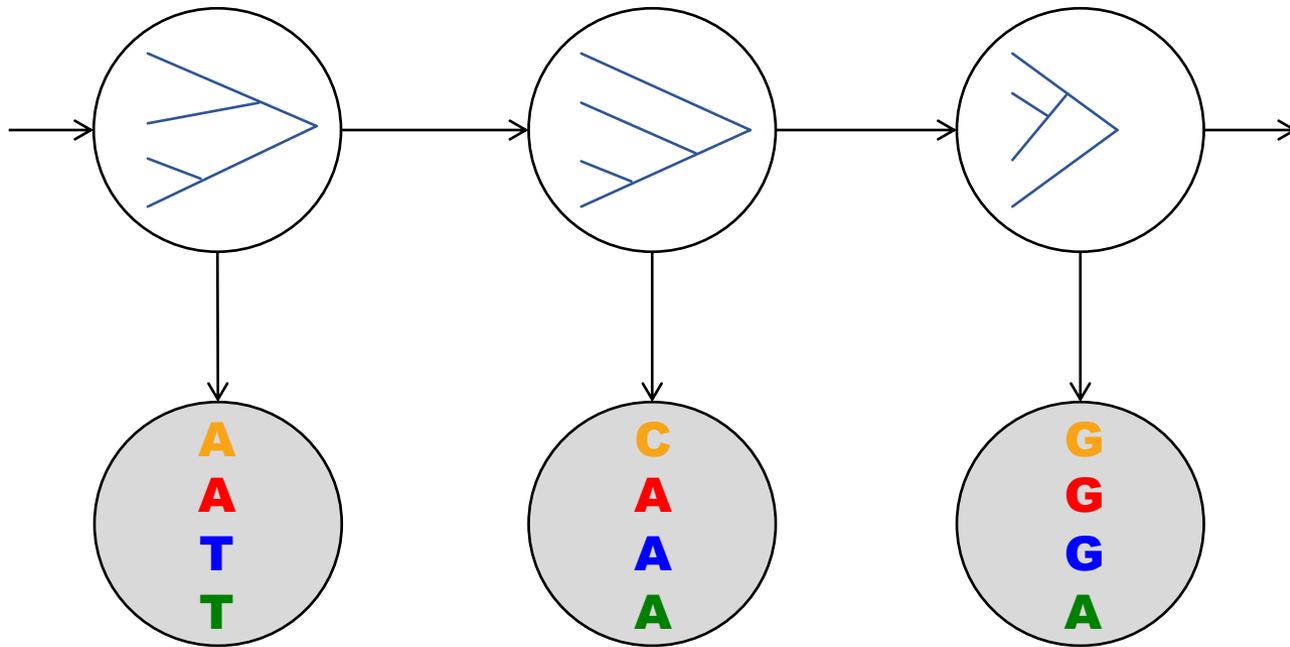
Tree changes along the genome!



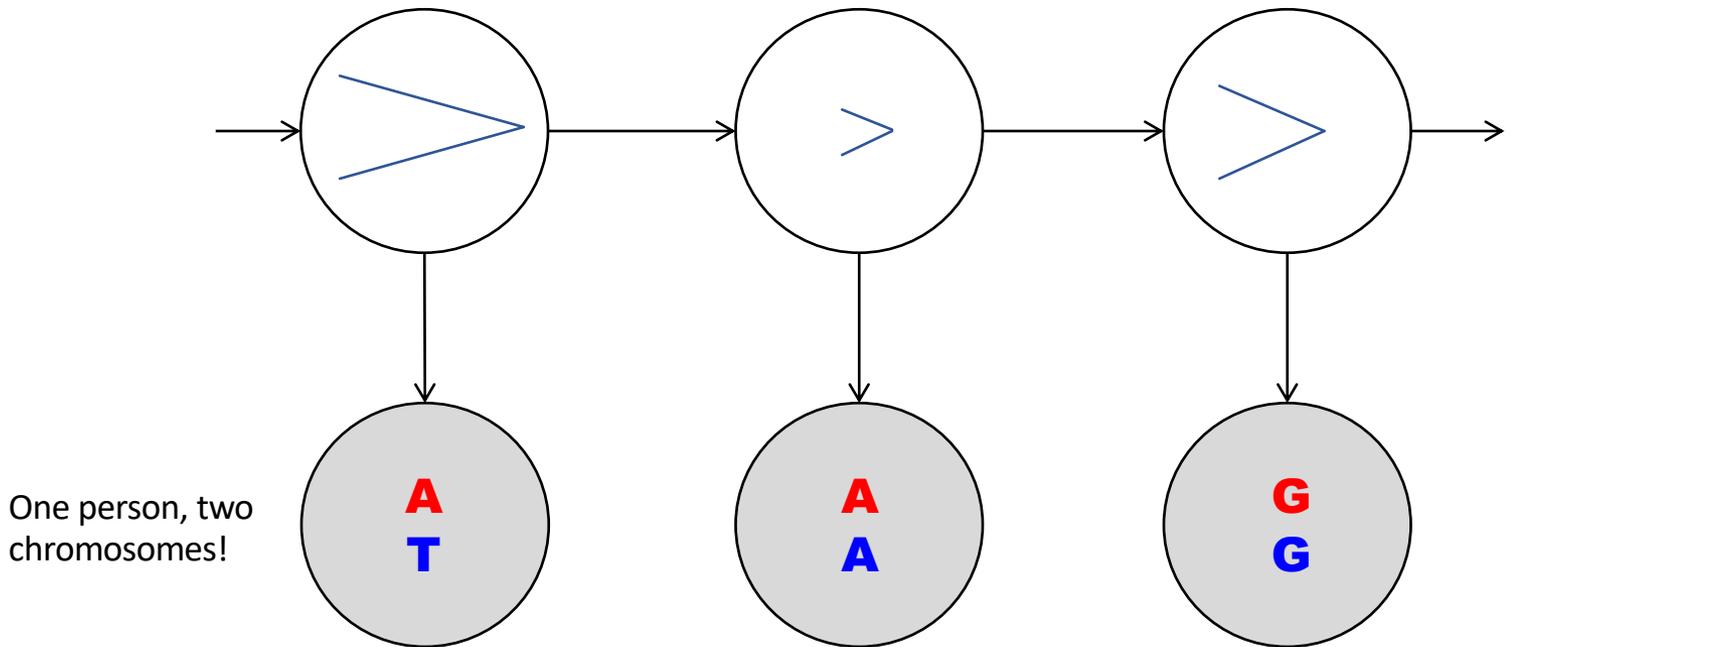
HMM observations: sequence data



HMM hidden states: the tree



Number of possible trees grows exponentially...



PSMC: pairwise sequentially Markovian coalescent

- The distribution of pairwise coalescence times should be **exponential** with parameter 1
- If this differs from the exponential distribution, there were probably **population size changes**
- If all coalescence times are very recent, small population size
- If all coalescence times are very ancient, large population size

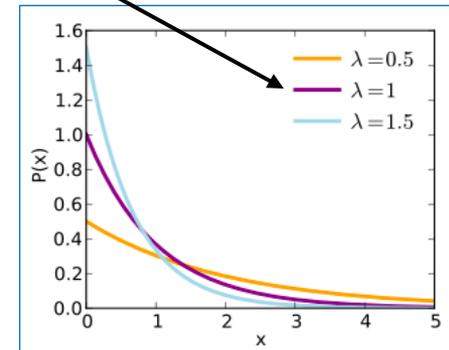
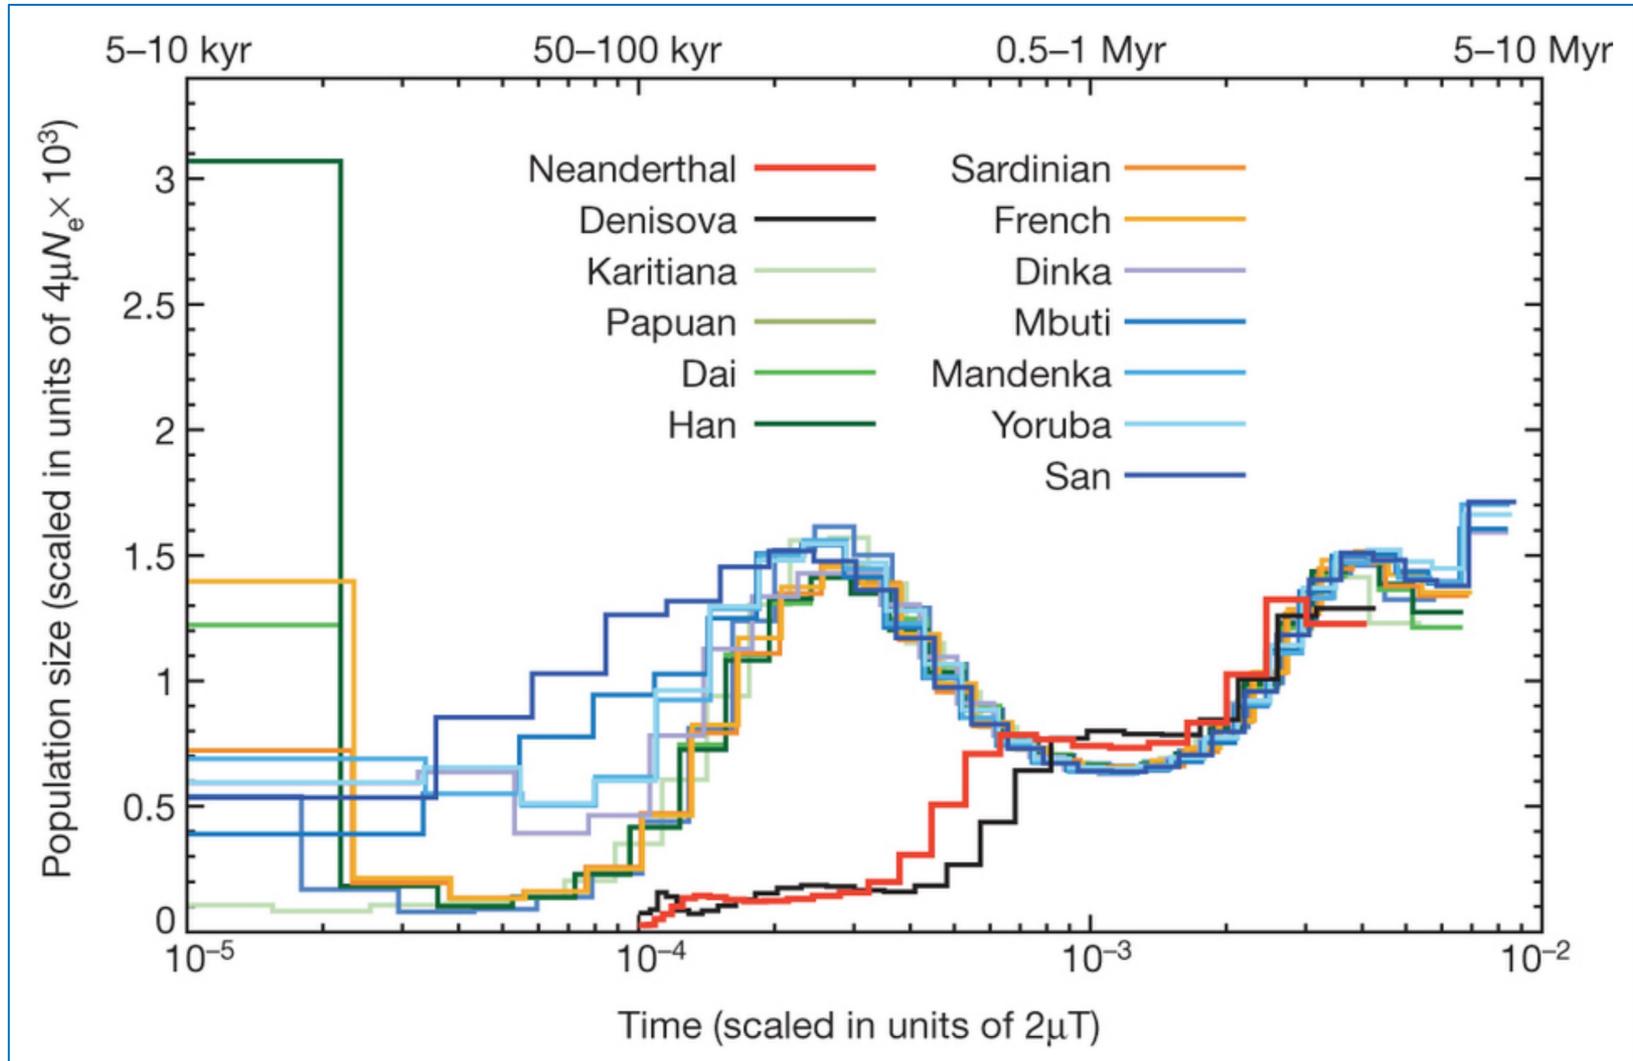


Image: wikipedia

PSMC: an HMM for two sequences



“The complete genome sequence of a Neanderthal from the Altai Mountains”, Prüfer et al (2014)

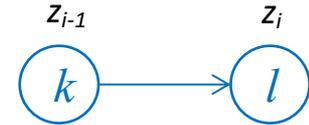
HMM definition

K = num hidden states, B = num emitted states

- Transition probabilities:

($K \times K$ matrix)

$$a_{kl} = P(z_i = l | z_{i-1} = k)$$



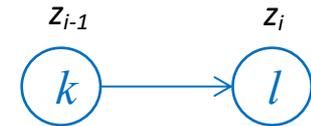
HMM definition

K = num hidden states, B = num emitted states

- Transition probabilities:

($K \times K$ matrix)

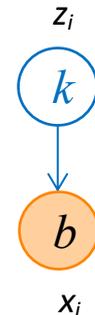
$$a_{kl} = P(z_i = l | z_{i-1} = k)$$



- Emission probabilities:

($K \times B$ matrix)

$$e_k(b) = P(x_i = b | z_i = k)$$



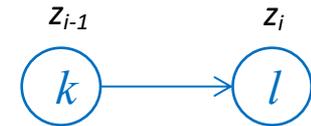
HMM definition

K = num hidden states, B = num emitted states

- Transition probabilities:

($K \times K$ matrix)

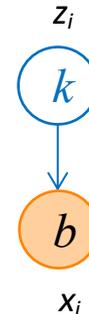
$$a_{kl} = P(z_i = l | z_{i-1} = k)$$



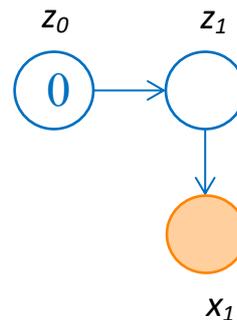
- Emission probabilities:

($K \times B$ matrix)

$$e_k(b) = P(x_i = b | z_i = k)$$



- A way to deal with the first state:



$$\pi_k = p(z_1 = k)$$

Z_0 = Special start state with no emission

π_k = probability of starting in state k

($K \times 1$ vector)

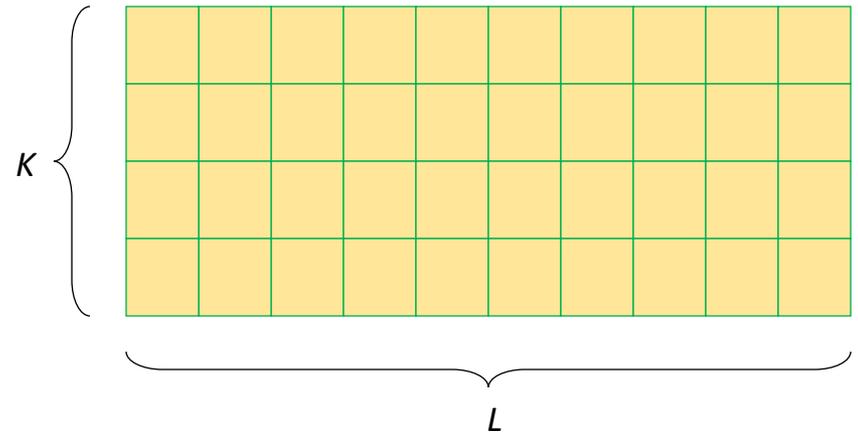
Do this way for lab!

Viterbi Algorithm

- **Input:** observed sequence (x_1, x_2, \dots, x_L) and transition/emission probabilities (\mathbf{a} and \mathbf{e} matrices)
- **Output:** most probable (i.e. most likely) hidden state sequence \mathbf{z}^*

Viterbi Algorithm

- **Input:** observed sequence (x_1, x_2, \dots, x_L) and transition/emission probabilities (\mathbf{a} and \mathbf{e} matrices)
- **Output:** most probable (i.e. most likely) hidden state sequence \mathbf{z}^*
- **Initialization:** create a $K \times L$ matrix, this will be our dynamic programming (DP) table V

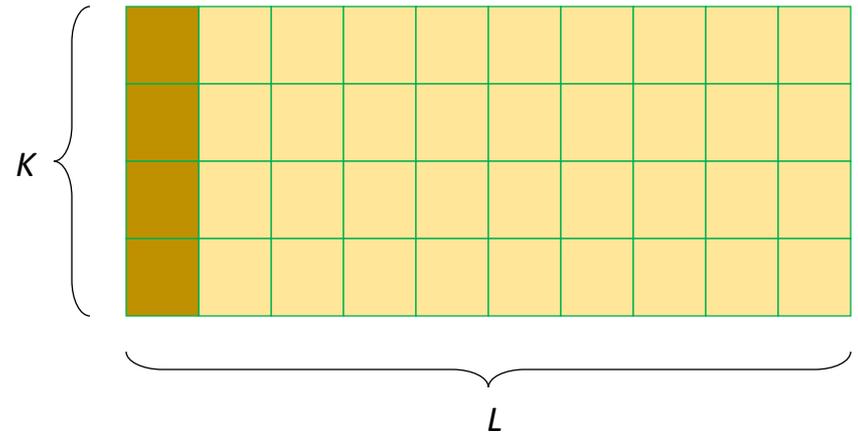


Viterbi Algorithm

- **Input:** observed sequence (x_1, x_2, \dots, x_L) and transition/emission probabilities (\mathbf{a} and \mathbf{e} matrices)
- **Output:** most probable (i.e. most likely) hidden state sequence \mathbf{z}^*
- **Initialization:** create a $K \times L$ matrix, this will be our dynamic programming (DP) table

$$V_k(1) = \pi_k e_k(x_1)$$

(Note: there are lots of ways to initialize, this avoids a special start state.)



Viterbi Algorithm

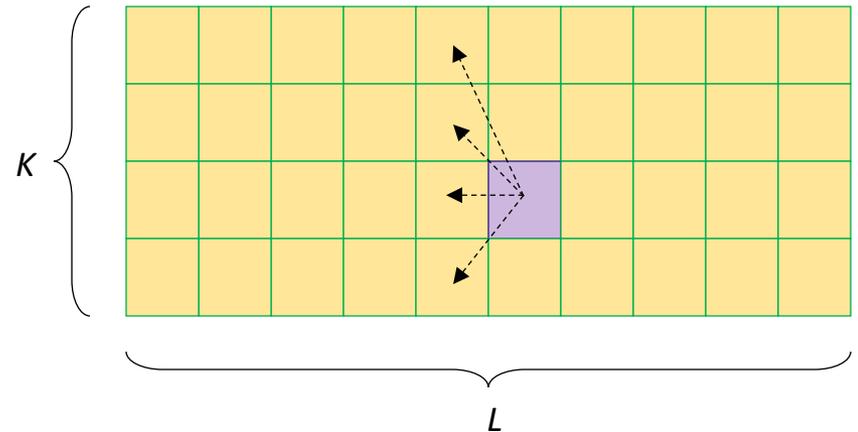
- **Input:** observed sequence (x_1, x_2, \dots, x_L) and transition/emission probabilities (\mathbf{a} and \mathbf{e} matrices)
- **Output:** most probable (i.e. most likely) hidden state sequence \mathbf{z}^*
- **Initialization:** create a $K \times L$ matrix, this will be our dynamic programming (DP) table

$$V_k(1) = \pi_k e_k(x_1)$$

(Note: there are lots of ways to initialize, this avoids a special start state.)

- **Recursion:**

$$V_k(i) = e_k(x_i) \cdot \max_l \{ V_l(i-1) \cdot a_{lk} \}$$



Viterbi Algorithm

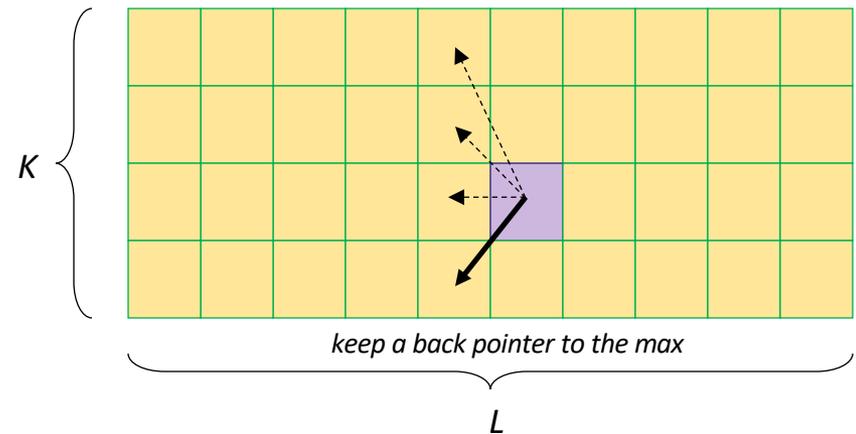
- **Input:** observed sequence (x_1, x_2, \dots, x_L) and transition/emission probabilities (\mathbf{a} and \mathbf{e} matrices)
- **Output:** most probable (i.e. most likely) hidden state sequence \mathbf{z}^*
- **Initialization:** create a $K \times L$ matrix, this will be our dynamic programming (DP) table

$$V_k(1) = \pi_k e_k(x_1)$$

(Note: there are lots of ways to initialize, this avoids a special start state.)

- **Recursion:**

$$V_k(i) = e_k(x_i) \cdot \max_l \{ V_l(i-1) \cdot a_{lk} \}$$



Viterbi Algorithm

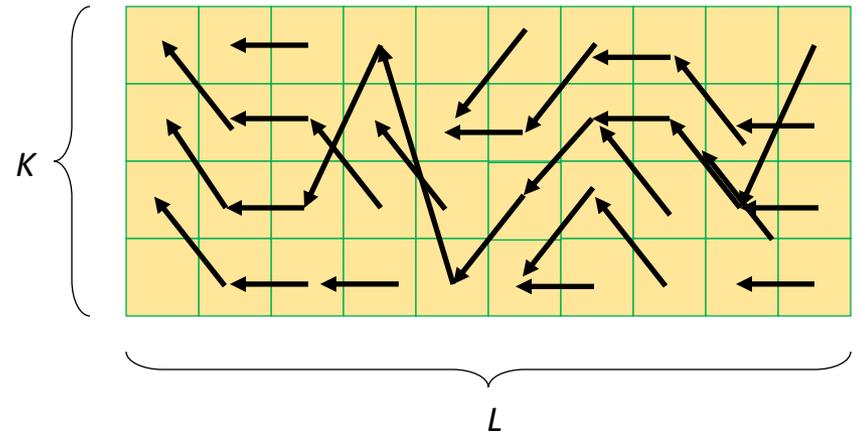
- **Input:** observed sequence (x_1, x_2, \dots, x_L) and transition/emission probabilities (\mathbf{a} and \mathbf{e} matrices)
- **Output:** most probable (i.e. most likely) hidden state sequence \mathbf{z}^*
- **Initialization:** create a $K \times L$ matrix, this will be our dynamic programming (DP) table

$$V_k(1) = \pi_k e_k(x_1)$$

(Note: there are lots of ways to initialize, this avoids a special start state.)

- **Recursion:**

$$V_k(i) = e_k(x_i) \cdot \max_l \{ V_l(i-1) \cdot a_{lk} \}$$



Viterbi Algorithm

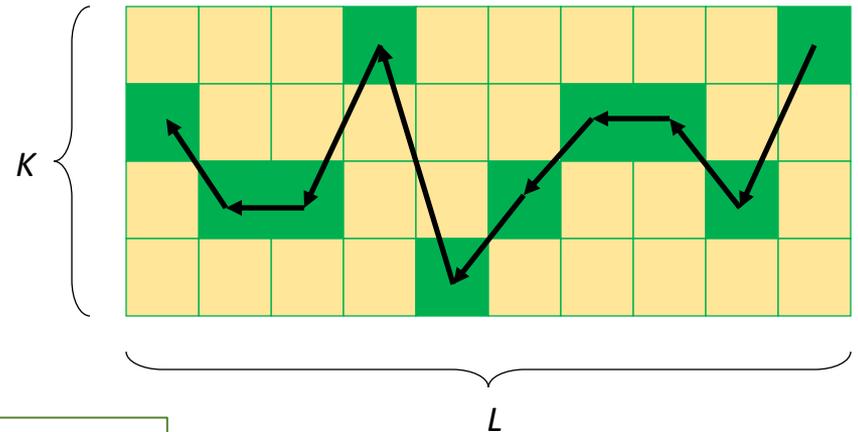
- **Input:** observed sequence (x_1, x_2, \dots, x_L) and transition/emission probabilities (\mathbf{a} and \mathbf{e} matrices)
- **Output:** most probable (i.e. most likely) hidden state sequence \mathbf{z}^*
- **Initialization:** create a $K \times L$ matrix, this will be our dynamic programming (DP) table

$$V_k(1) = \pi_k e_k(x_1)$$

(Note: there are lots of ways to initialize, this avoids a special start state.)

- **Recursion:**

$$V_k(i) = e_k(x_i) \cdot \max_l \{ V_l(i-1) \cdot a_{lk} \}$$



- **Termination and traceback:**

$$P(\vec{x}, \vec{z}^*) = \max_k \{ V_k(L) \}$$

Viterbi Algorithm

- **Input:** observed sequence (x_1, x_2, \dots, x_L) and transition/emission probabilities (\mathbf{a} and \mathbf{e} matrices)
- **Output:** most probable (i.e. most likely) hidden state sequence \mathbf{z}^*
- **Initialization:** create a $K \times L$ matrix, this will be our dynamic programming (DP) table

$$V_k(1) = \pi_k e_k(x_1)$$

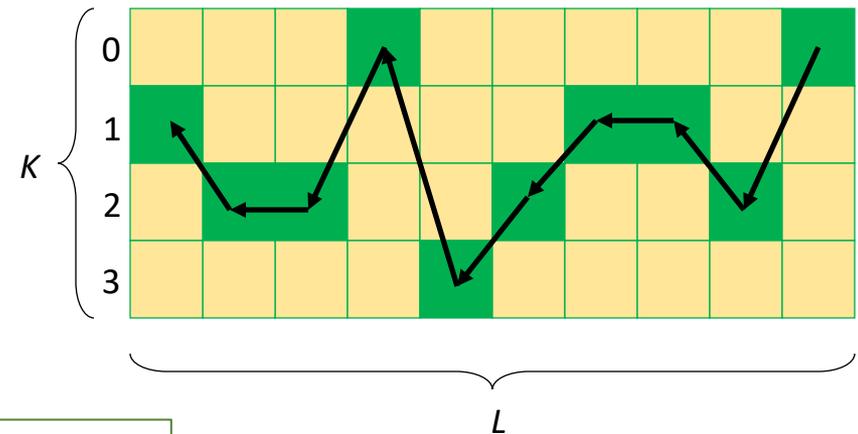
(Note: there are lots of ways to initialize, this avoids a special start state.)

- **Recursion:**

$$V_k(i) = e_k(x_i) \cdot \max_l \{ V_l(i-1) \cdot a_{lk} \}$$

- **Termination and traceback:**

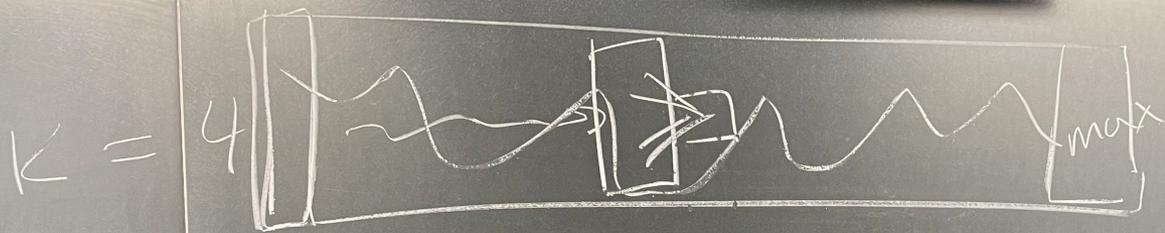
$$P(\vec{x}, \vec{z}^*) = \max_k \{ V_k(L) \}$$



$$\mathbf{z}^* = (1, 2, 2, 0, 3, 2, 1, 1, 2, 0)$$

K states, L observations: what is the complexity of the Viterbi algorithm?

Viterbi algorithm worksheet



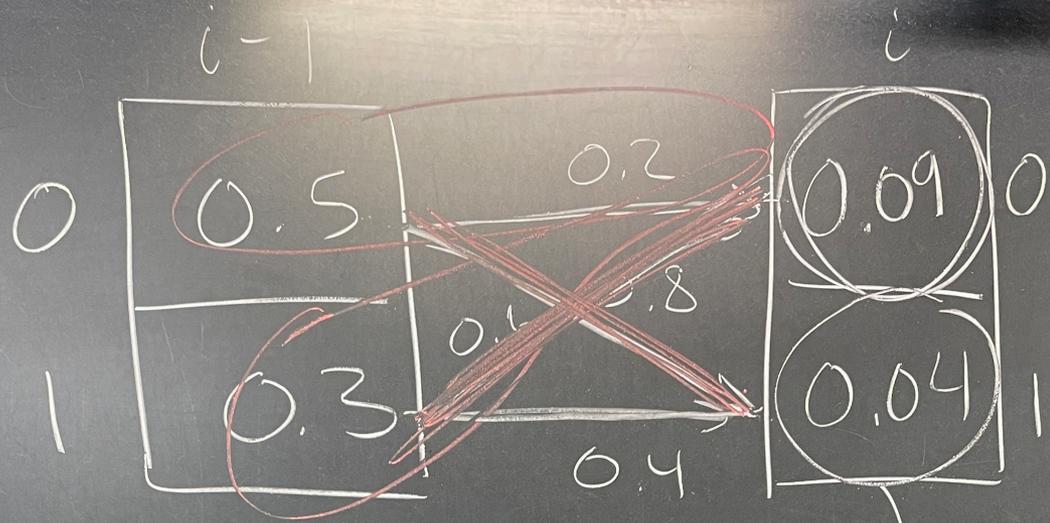
$$L = 100,000$$

① $K \cdot K \cdot K \dots K = K^L$
 $i=1 \quad i=2 \quad i=3 \quad i=L$

② at least ~~$O(KL)$~~

$$O(LK^2)$$

3



State seq



$$[z_1, z_2] = \mathbf{z}^*$$

$$\begin{aligned}
 V_0(i) &= e_0(A) \cdot \max \{ V_0(i-1)a_{00}, V_1(i-1)a_{10} \} \\
 &= 0.5 \cdot \max \{ \underbrace{0.5 \cdot 0.2}_{0.1}, \underbrace{0.3 \cdot 0.6}_{0.18} \} \\
 &= \boxed{0.09}
 \end{aligned}$$

$$\begin{aligned}
 V_1(i) &= 0.1 \cdot \max \{ \underbrace{0.5 \cdot 0.8}_{0.4}, 0.3 \cdot 0.4 \} \\
 &= 0.1 \cdot 0.4 \\
 &= \boxed{0.04}
 \end{aligned}$$

⑤ log space!

④ $i+1$

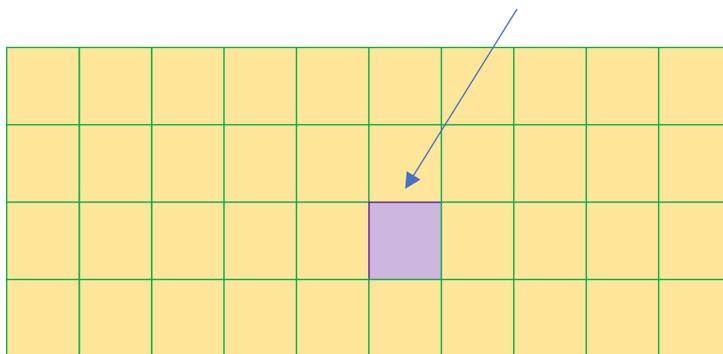
0.0024
0.0216

state seq
101

Forward-backward Algorithm

- **Input:** observed sequence $x=(x_1,x_2,\dots,x_L)$ and transition/emission probabilities (\mathbf{a} and \mathbf{e} matrices)
- **Output:** posterior probability of being in each hidden state at each time point $P(Z_i=k|x)$

What is the probability of being in this state, conditional on the observations?

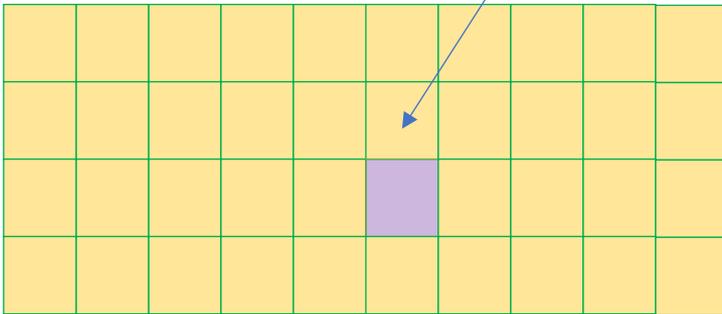


Want to compute $P(Z_i=k|x_1,x_2,\dots,x_L) = P(Z_i=k, x_1,x_2,\dots,x_L) / P(x_1,x_2,\dots,x_L)$

Forward-backward Algorithm

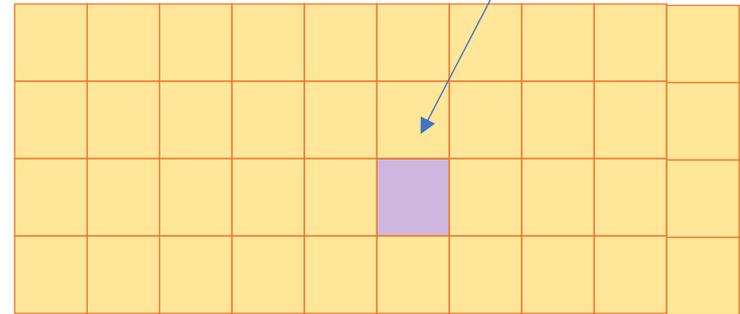
Forward matrix $f_k(i)$

What is the probability of being in this state, conditional on all the earlier observations



Backward matrix $b_k(i)$

What is the probability of all the later observations, conditional on being in this state



Want to compute $P(Z_i=k | x_1, x_2, \dots, x_L) = P(Z_i=k, x_1, x_2, \dots, x_L) / P(x_1, x_2, \dots, x_L)$

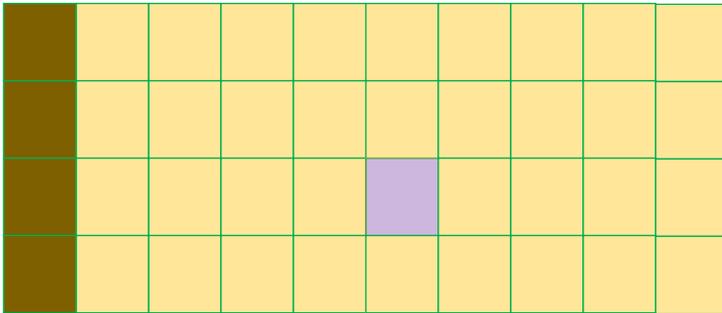
But we can break this up: $P(Z_i=k, x_1, x_2, \dots, x_L) = P(Z_i=k, x_1, x_2, \dots, x_i) P(x_{i+1}, x_2, \dots, x_L | Z_i=k)$

Probability of being in state k at time i , given the sequence up to time i : Forward algorithm

Probability of the sequence after time i , given that you are in state k at time i : Backward algorithm

Forward Algorithm

1. $f_k(1) = \pi_k e_k(x_1)$

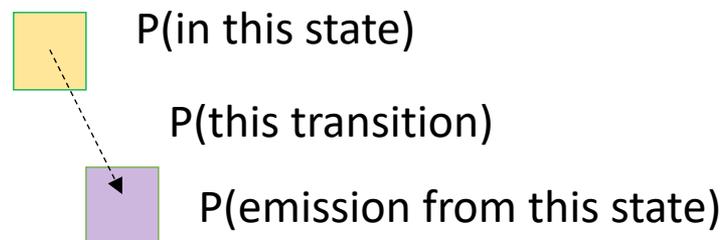
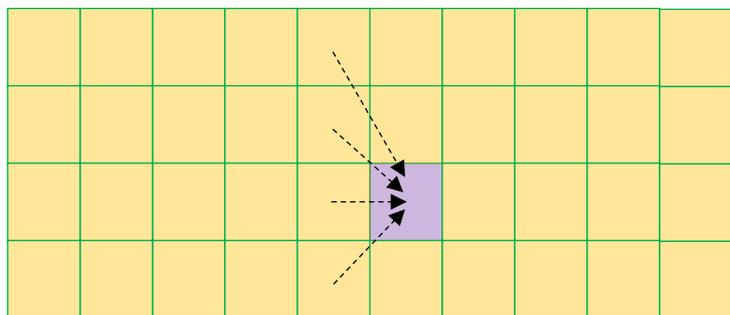


$$P(Z_i=k, x_1, x_2 \dots x_i)$$

Probability of being in state k at time i , given the sequence up to time i : Forward algorithm

Forward Algorithm

1. $f_k(1) = \pi_k e_k(x_1)$
2. $f_k(i) = \sum_{j=1}^K f_j(i-1) a_{jk} e_k(x_i)$



$$P(Z_i=k, x_1, x_2 \dots x_i)$$

Probability of being in state k at time i , given the sequence up to time i : Forward algorithm

Forward Backward Algorithm

Goal
 (posterior)

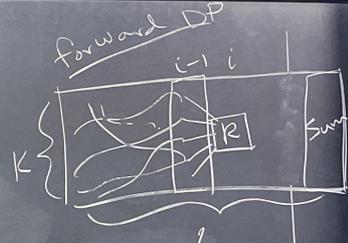
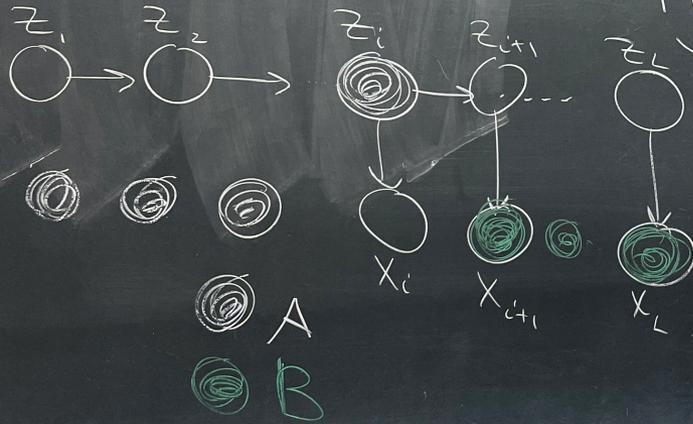
$$P(z_i=k | \vec{x}) = \frac{P(\vec{x}, z_i=k)}{P(\vec{x})}$$

up to i
for now

$$P(\vec{x}, z_i=k) = P(x_1, x_2, \dots, x_i, z_i=k, x_{i+1}, x_{i+2}, \dots, x_L)$$

BAYES

$$= P(x_1, x_2, \dots, x_i, z_i=k) \cdot P(x_{i+1}, \dots, x_L | z_i=k, x_1, x_2, \dots, x_i)$$



forward probabilities

backward probabilities

Forward Algorithm

$f_k(i)$ = prob of observing x_1, \dots, x_i
+ ending in state k

$$= P(x_1, \dots, x_i, z_i = k)$$

all possible prev states

recursion: $f_k(i) = e_k(x_i) \sum_l f_l(i-1) a_{lk}$

initialization (Same as viterbi)
 integrating over all possibilities

$$f_k(1) = \pi_k e_k(x_1)$$

termination $P(\vec{x}) = \sum_k f_k(L)$

together

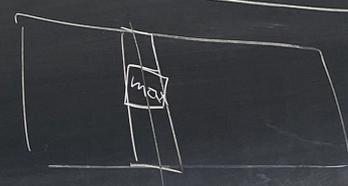
$$p(\vec{x}, z_i = k) = f_k(i) b_k(i)$$

$$p(z_i = k | \vec{x}) = \frac{f_k(i) b_k(i)}{P(\vec{x})}$$

posterior prob

posterior decoding

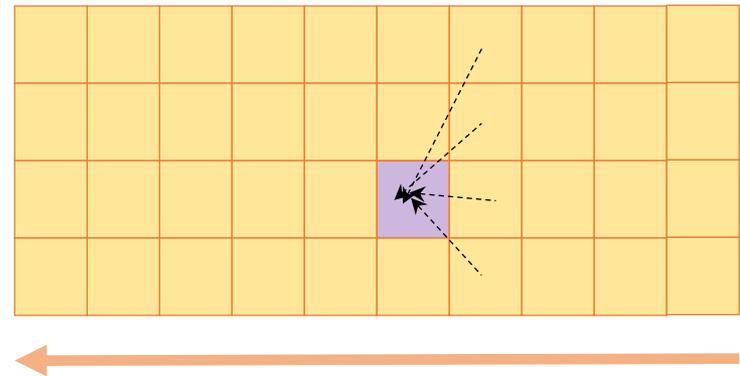
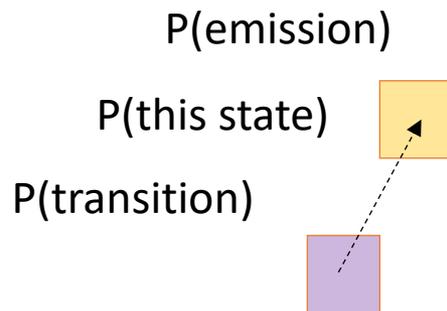
$$\hat{z}_i = \operatorname{argmax}_k P(z_i = k | \vec{x})$$



Backward Algorithm

1 $b_k(K) = 1$

2 $b_k(i) = \sum_{j=1}^K b_j(i+1) a_{ij} e_j(x_{i+1})$



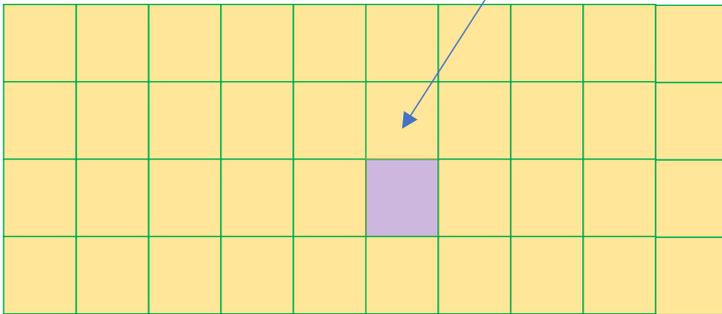
$$P(x_{i+1}, x_2 \dots x_L | Z_i = k)$$

Probability of the sequence after time i, given that you are in state k at time i: Backward algorithm

Forward-backward Algorithm

Forward matrix $f_k(i)$

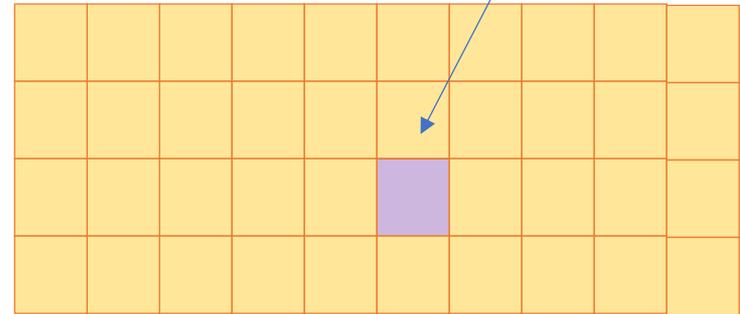
What is the probability of being in this state, conditional on all the earlier observations



×

Backward matrix $b_k(i)$

What is the probability of all the later observations, conditional on being in this state



Want to compute $P(Z_i=k | x_1, x_2, \dots, x_L) = P(Z_i=k, x_1, x_2, \dots, x_L) / P(x_1, x_2, \dots, x_L)$

But we can break this up: $P(Z_i=k, x_1, x_2, \dots, x_L) = P(Z_i=k, x_1, x_2, \dots, x_i) P(x_{i+1}, x_2, \dots, x_L | Z_i=k)$

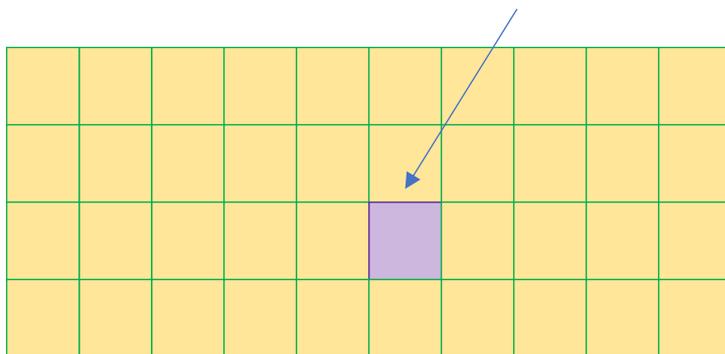
Probability of being in state k at time i , given the sequence up to time i : Forward algorithm

Probability of the sequence after time i , given that you are in state k at time i : Backward algorithm

Forward-backward Algorithm

- **Input:** observed sequence $x=(x_1,x_2,\dots,x_L)$ and transition/emission probabilities (\mathbf{a} and \mathbf{e} matrices)
- **Output:** posterior probability of being in each hidden state at each time point $P(Z_i=k|x)$

What is the probability of being in this state, conditional on the observations?



Want to compute $P(Z_i=k|x_1,x_2,\dots,x_L) = P(Z_i=k, x_1,x_2,\dots,x_L) / P(x_1,x_2,\dots,x_L)$

What is the runtime of the forward-backward algorithm?

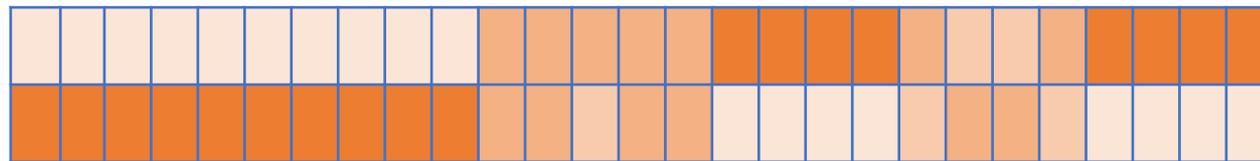
Parameter estimation

So far we assumed that we know the emission and transition probabilities, and the initial state distribution. If we don't know these, can we estimate them?

Yes – using the Baum-Welch algorithm; a special case of the Expectation-Maximization (EM) algorithm:

0. Guess some parameters and run the F-B algorithm

1. Use estimated parameters to run the F-B algorithm

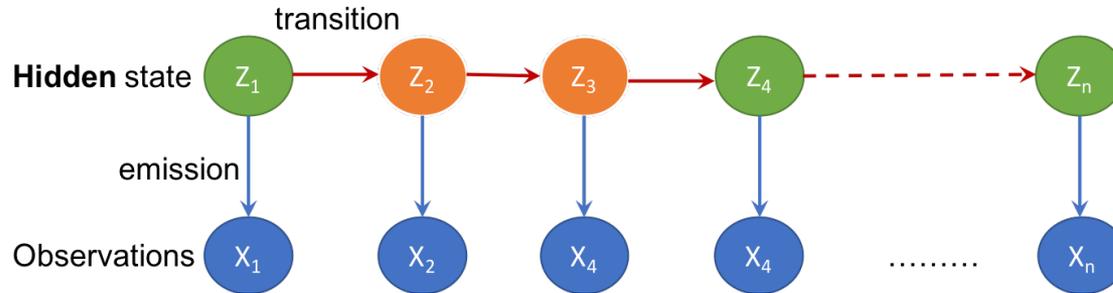


2. Use the output of the F-B algorithm to estimate the parameters

(estimate the parameters by using the empirical probabilities from the FB matrix)

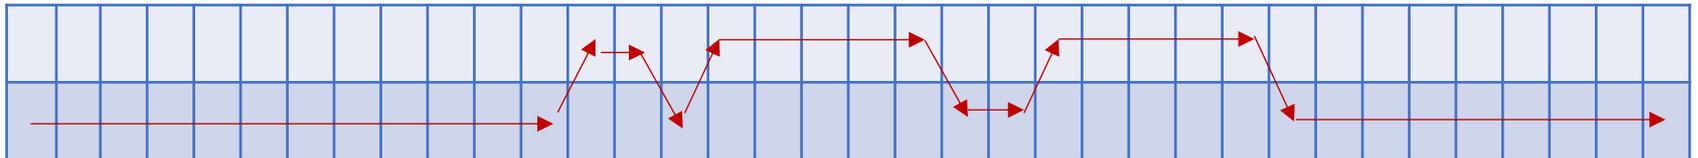
Summary of HMMs

Structure of a Hidden Markov Model: Observations, hidden states, emissions, transitions

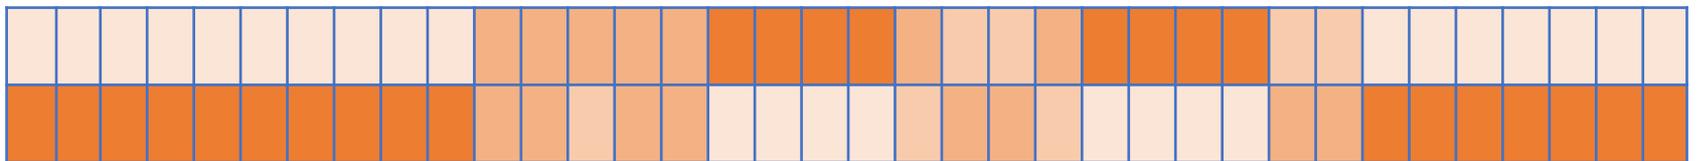


Three inference problems:

1) What is the most likely sequence of hidden states?: Viterbi algorithm



2) What is the most likely hidden state at each observation? Forward-Backward algorithm



3) What are the parameters? Baum-Welch algorithm