



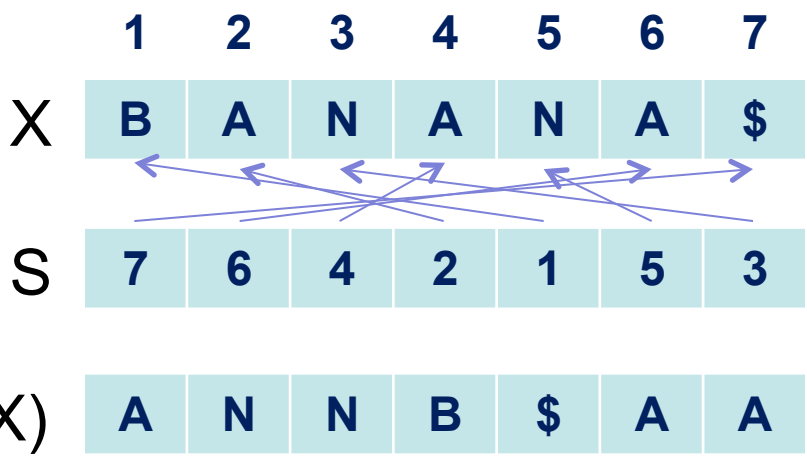
Review: Suffix Arrays and BWT

\$BANANA	1	\$BANANA
A\$BANAN	2	A\$BANAN
ANA\$BAN	3	ANA\$BAN
ANANA\$B	4	ANANA\$B
BANANA\$	5	BANANA\$
NA\$BANA	6	NA\$BANA
NANA\$BA	7	NANA\$BA

Suffixes are sorted in the BWT matrix

Define suffix array S :

$S(i) = j$, where $X_j \dots X_n$ is the i -th suffix lexicographically



BWT(X) constructed from S:
At each position, take the letter to the left of the one pointed by S



Review: Reconstructing BANANA

\$BANANA
A\$BANAN
ANA\$BAN
ANANA\$B
BANANA\$
NA\$BANA
NANA\$BA

BWT matrix of string 'BANANA'

A N N B \$ A A

C()	1	5	5	4	0	1	1
index i	1	1	2	1	1	2	3
LF()	2	6	7	5	1	3	4

C(a) character array:
letter occs before a

i: indicating i-th occur.
of 'a' in BWT

$LF() = C() + i$

```

Reconstruct BANANA:
S := ""; r := 1; c := BWT[r];
UNTIL c = '$' {
    S := cS;
    r := LF(r);
    c := BWT(r); }

```



Searching for query “ANA”

```
$BANANA
A$BANAN
ANA$BAN
ANANA$B
BANANA$
NA$BANA
NANA$BA
```

BWT matrix of string ‘BANANA’

```
Let
LFC(r, a) = C(a) + i, where i = #'a's up to r in BWT

ExactMatch(W[1...k]) {
    a := W[k];
    low := C(a) + 1;
    high := C(a+1); // a+1: lexicographically next char
    i := k - 1;
    while (low <= high && i >= 1) {
        a = W[i];
        low = LFC(low - 1, a) + 1;
        high = LFC(high, a);
        i := i - 1; }
    return (low, high);
}
```



BWT Index Construction

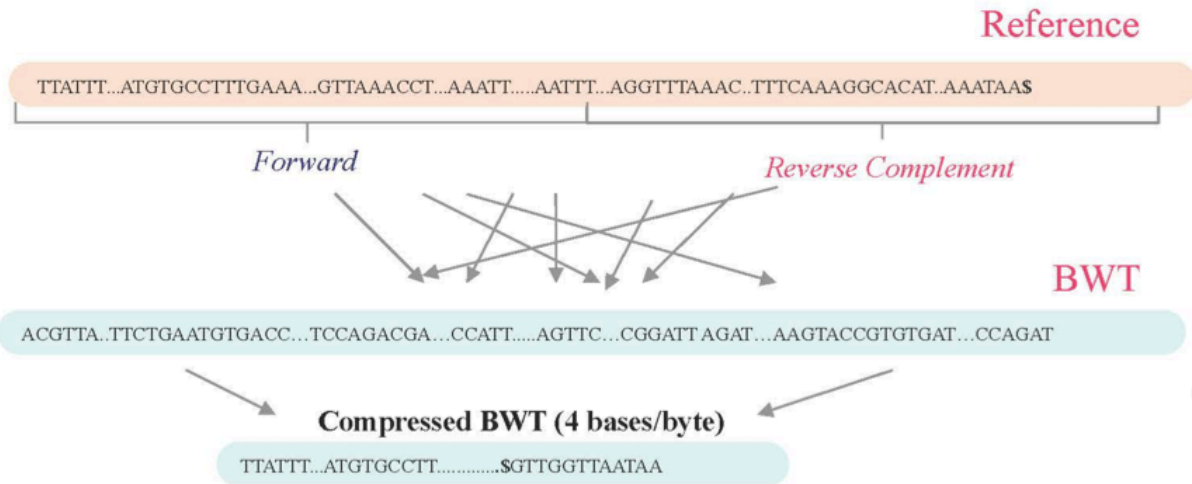


Reference Sequence Construction

BWT Construction

BWT-auxiliary Structure Construction (C & O arrays) and Compression

.bwt



C-array

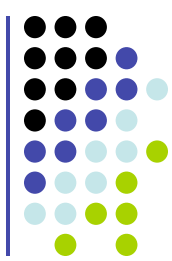
S:	0
T:	55000
C:	1044814
G:	7814189
A:	1

O-array

	T	G	C	A
0:	1	0	0	0
1:	2	0	0	0
2:
3:
..
G -1:

SA

0:	G -1
1:	64
2:	144814
3:	781414689
..	...
G -1:	1484



BWA Inexact Match

Allow up to n mismatches/gaps

Backward search:

Given read W , keep track of multiple partial alignments

Partial alignment: (i, z, L, U)

i : current position

z : remaining non-matches allowed

L : current low

U : current high

```
 $I \leftarrow \emptyset$   
 $I \leftarrow I \cup \text{INEXRECUR}(W, i-1, z-1, k, l)$   
for each  $b \in \{A, C, G, T\}$  do  
   $k \leftarrow C(b) + O(b, k-1) + 1$   
   $l \leftarrow C(b) + O(b, l)$   
  if  $k \leq l$  then  
     $I \leftarrow I \cup \text{INEXRECUR}(W, i, z-1, k, l)$   
    if  $b = W[i]$  then  
       $I \leftarrow I \cup \text{INEXRECUR}(W, i-1, z, k, l)$   
    else  
       $I \leftarrow I \cup \text{INEXRECUR}(W, i-1, z-1, k, l)$ 
```



BWA Inexact Match

W = ACTGTGT
←

Partial alignment 4-tuple: (i = 4, z = 3, L, U)

Recursive step:

A	C	T	G	gap-ref	gap-read
AGT	CGT	TGT	GGT	∓GT	*GT
z-1	z-1	z	z-1	z-1	z-1
i-1	i-1	i-1	i-1	i-1	i
L ^A U ^A	L ^C U ^C	L ^T U ^T	L ^G U ^G	LU	L ^A U ^A L ^C U ^C L ^T U ^T L ^G U ^G
...GAGT	...GCGT	...GTGT	...GGGT	...G-GT	...GT[A/C/T/G]GT
...GTGT	...GTGT	...GTGT	...GTGT	...GTGT	...GT - GT

$$L^A = C(A) + O(A, L-1) + 1$$

$$U^A = C(A) + O(A, L)$$

```

I ← ∅
I ← I ∪ INEXRECUR(W, i-1, z-1, k, l)
for each b ∈ {A, C, G, T} do
  k ← C(b) + O(b, k-1) + 1
  l ← C(b) + O(b, l)
  if k ≤ l then
    I ← I ∪ INEXRECUR(W, i, z-1, k, l)
    if b = W[i] then
      I ← I ∪ INEXRECUR(W, i-1, z, k, l)
    else
      I ← I ∪ INEXRECUR(W, i-1, z-1, k, l)
  
```

BWA Heuristics

- Lower bound array D , where $D(i) :=$ **LB on number of differences** of exactly matching $R[0,i]$ with the reference (can be computed in $O(|R|)$ time \rightarrow check $n < D(i)$ instead of $n < 0$)
- Process best partial alignments first: use a *min*-priority **heap** to store alignment entries (instead of recursion)
- Prune out alignments considered sub-optimal (although they might have fewer than n differences):
dynamically adjust search parameters (e.g. n):
 - (1) stop if # top hits exceeds a threshold (=30),
 - (2) set $n = n_{best} + 1$, where n_{best} is the # of differences in top hit
- Seeding: limit the number of differences in the *seed* sequence (first k bp)
- Disallow indels at the ends of the read

Li H, Durbin R.

[Fast and accurate short read alignment with Burrows-Wheeler transform. Bioinformatics, 2009.](#)

7154 cites

Langmead B, Salzberg SL.

[Fast gapped-read alignment with Bowtie2. Nature Methods, 2012.](#)

3017 cites

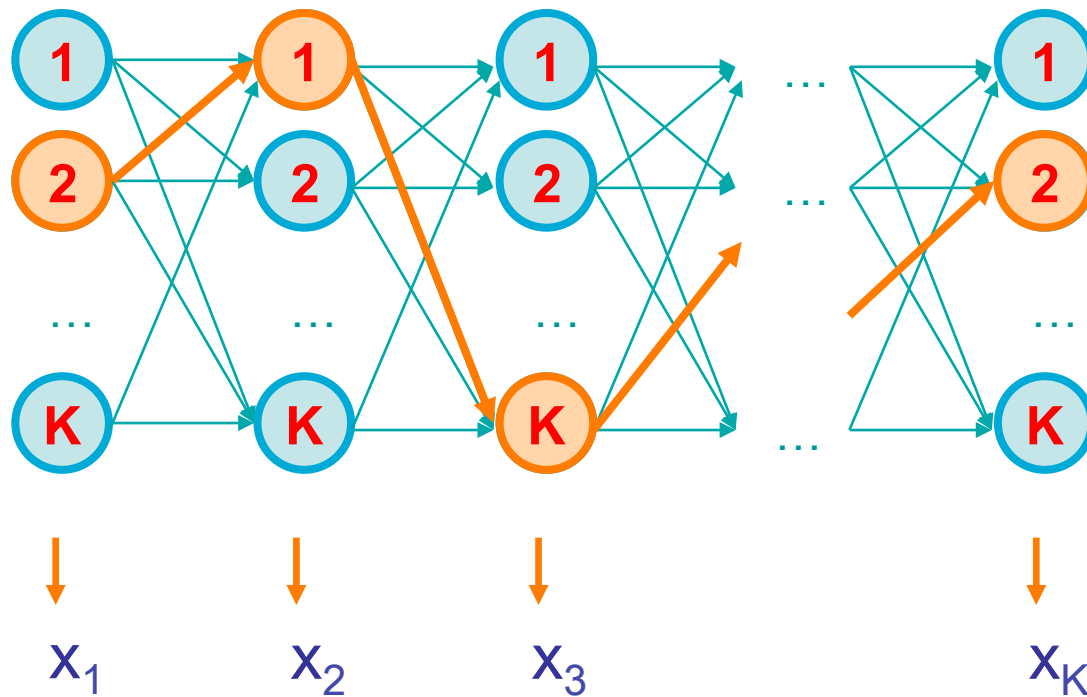
Li H

[Aligning sequence reads, clone sequences and assembly contigs with BWA-MEM](#)

Credit: Victoria Popic



Hidden Markov Models





Example: The Dishonest Casino

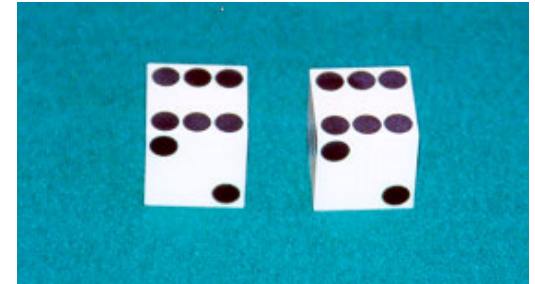
A casino has two dice:

- Fair die
 $P(1) = P(2) = P(3) = P(5) = P(6) = 1/6$
- Loaded die
 $P(1) = P(2) = P(3) = P(5) = 1/10$
 $P(6) = 1/2$

Casino player switches back-&-forth between fair and loaded die once every 20 turns

Game:

1. You bet \$1
2. You roll (always with a fair die)
3. Casino player rolls (maybe with fair die, maybe with loaded die)
4. Highest number wins \$2





Question # 1 – Evaluation

GIVEN

A sequence of rolls by the casino player

1245526462146146136136661664661636616366163616515615115146123562344

Prob = 1.3×10^{-35}

QUESTION

How likely is this sequence, given our model of how the casino works?

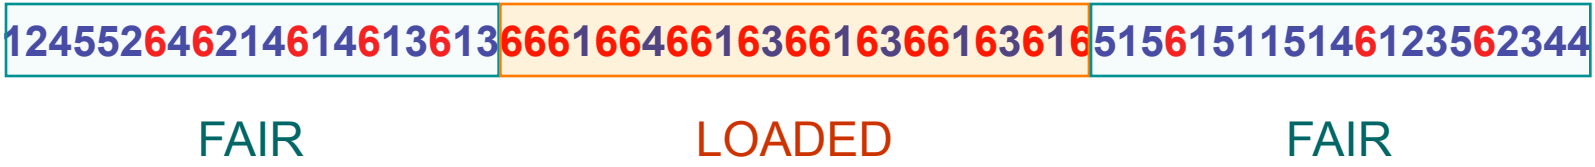
This is the **EVALUATION** problem in HMMs



Question # 2 – Decoding

GIVEN

A sequence of rolls by the casino player



QUESTION

What portion of the sequence was generated with the fair die, and what portion with the loaded die?

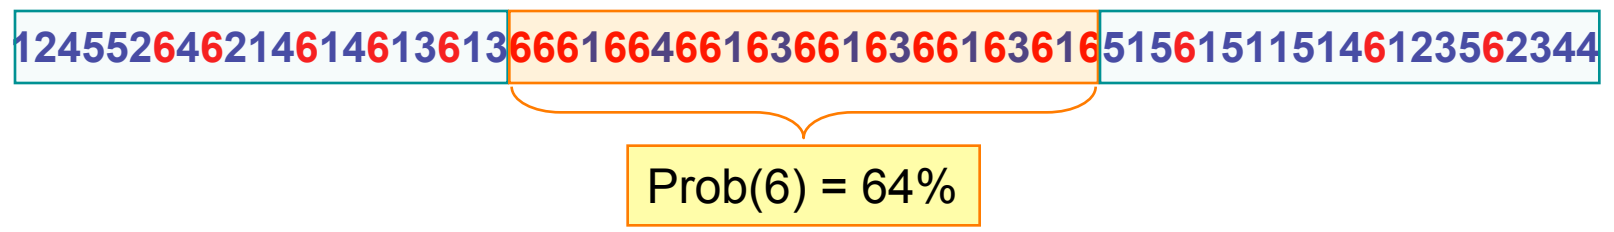
This is the **DECODING** question in HMMs



Question # 3 – Learning

GIVEN

A sequence of rolls by the casino player



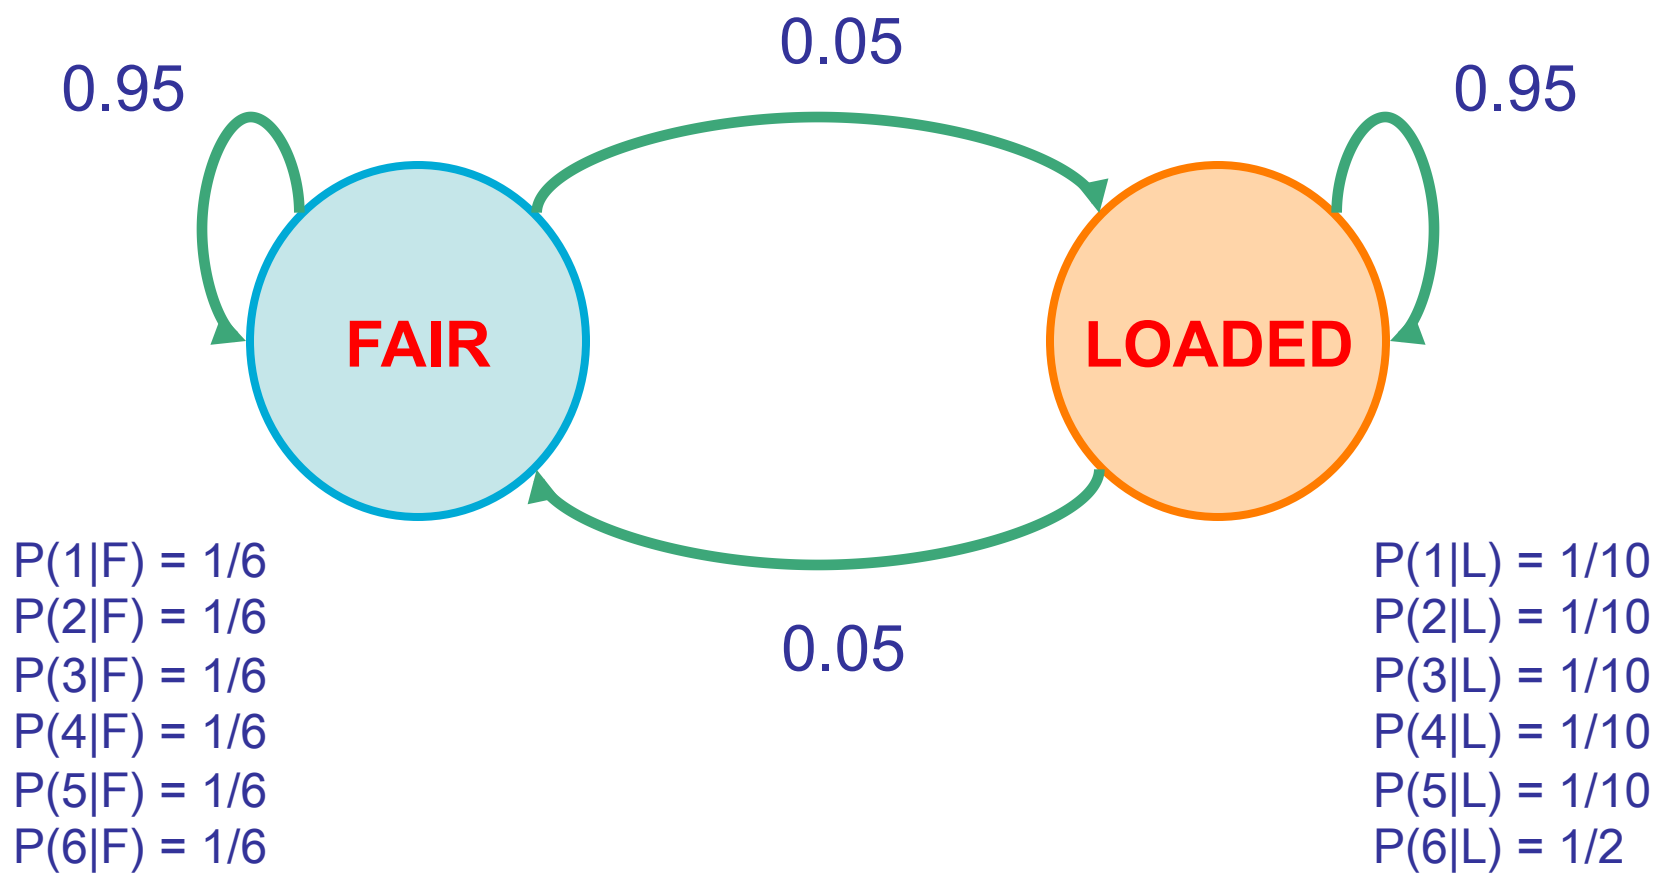
QUESTION

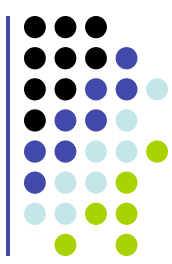
How “loaded” is the loaded die? How “fair” is the fair die? How often does the casino player change from fair to loaded, and back?

This is the **LEARNING** question in HMMs



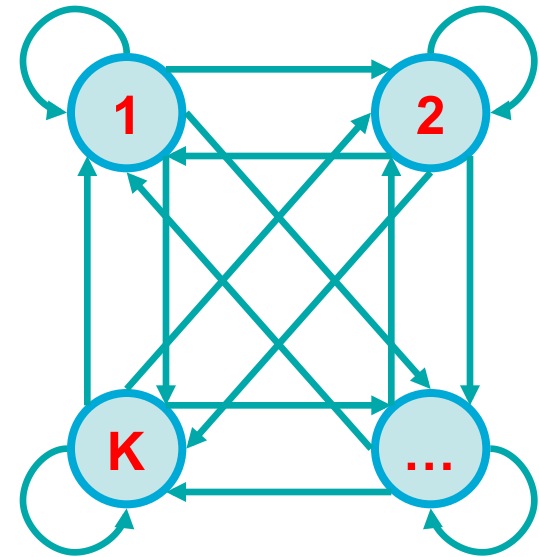
The dishonest casino model





A HMM is memory-less

At each time step t ,
the only thing that affects future states
is the current state π_t





Definition of a hidden Markov model

Definition: A hidden Markov model (HMM)

- **Alphabet** $\Sigma = \{ b_1, b_2, \dots, b_M \}$
- **Set of states** $Q = \{ 1, \dots, K \}$
- **Transition probabilities** between any two states

a_{ij} = transition prob from state i to state j
 $a_{i1} + \dots + a_{iK} = 1$, for all states $i = 1 \dots K$

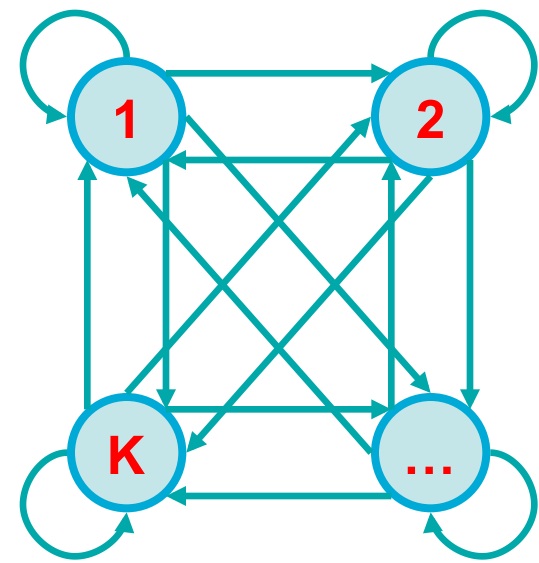
- **Start probabilities** a_{0i}

$$a_{01} + \dots + a_{0K} = 1$$

~~End Probabilities a_{i0}~~
in Durbin; not needed

- **Emission probabilities** within each state

$$e_i(b) = P(x_i = b \mid \pi_i = k)$$
$$e_i(b_1) + \dots + e_i(b_M) = 1, \text{ for all states } i = 1 \dots K$$

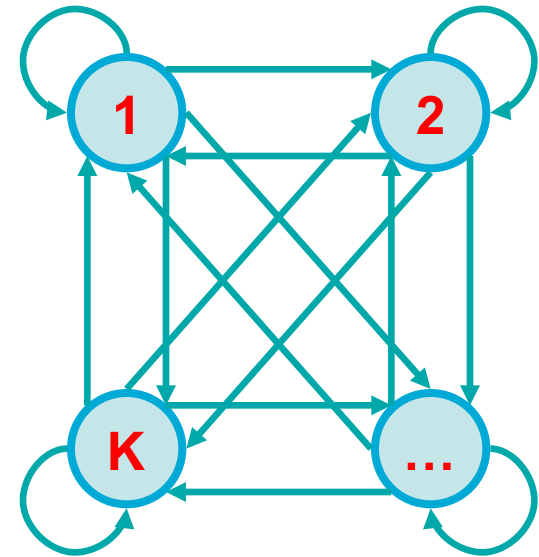




A HMM is memory-less

At each time step t ,
the only thing that affects future states
is the current state π_t

$$\begin{aligned} P(\pi_{t+1} = k \mid \text{“whatever happened so far”}) &= \\ P(\pi_{t+1} = k \mid \pi_1, \pi_2, \dots, \pi_t, x_1, x_2, \dots, x_t) &= \\ P(\pi_{t+1} = k \mid \pi_t) \end{aligned}$$

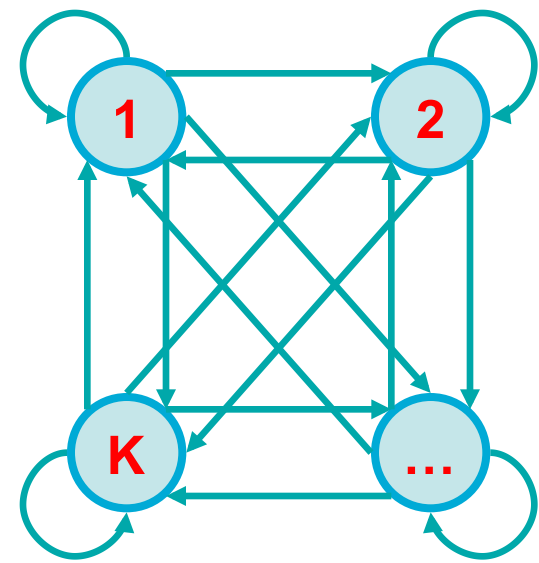




A HMM is memory-less

At each time step t ,
the only thing that affects x_t
is the current state π_t

$$\begin{aligned} P(x_t = b \mid \text{“whatever happened so far”}) &= \\ P(x_t = b \mid \pi_1, \pi_2, \dots, \pi_t, x_1, x_2, \dots, x_{t-1}) &= \\ P(x_t = b \mid \pi_t) \end{aligned}$$

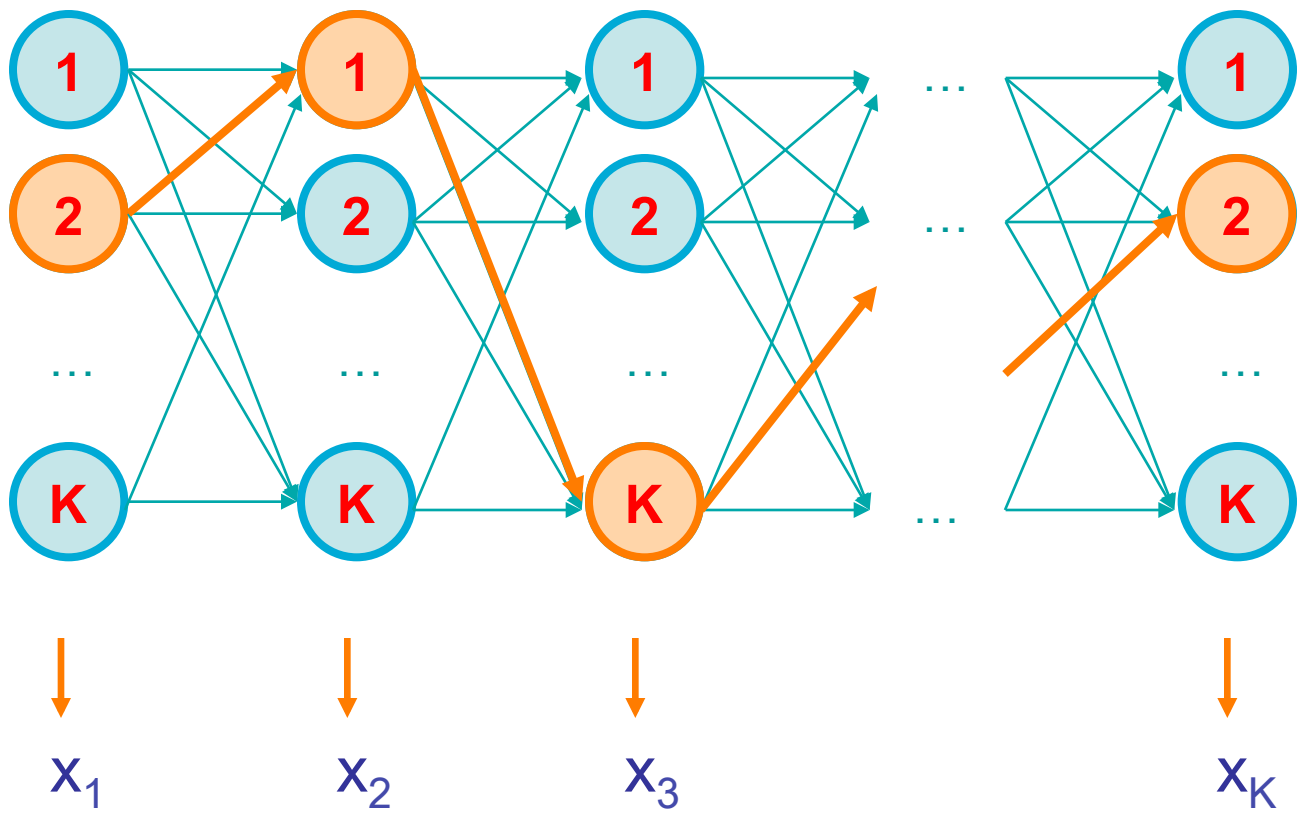




A parse of a sequence

Given a sequence $x = x_1 \dots x_N$,

A parse of x is a sequence of states $\pi = \pi_1, \dots, \pi_N$

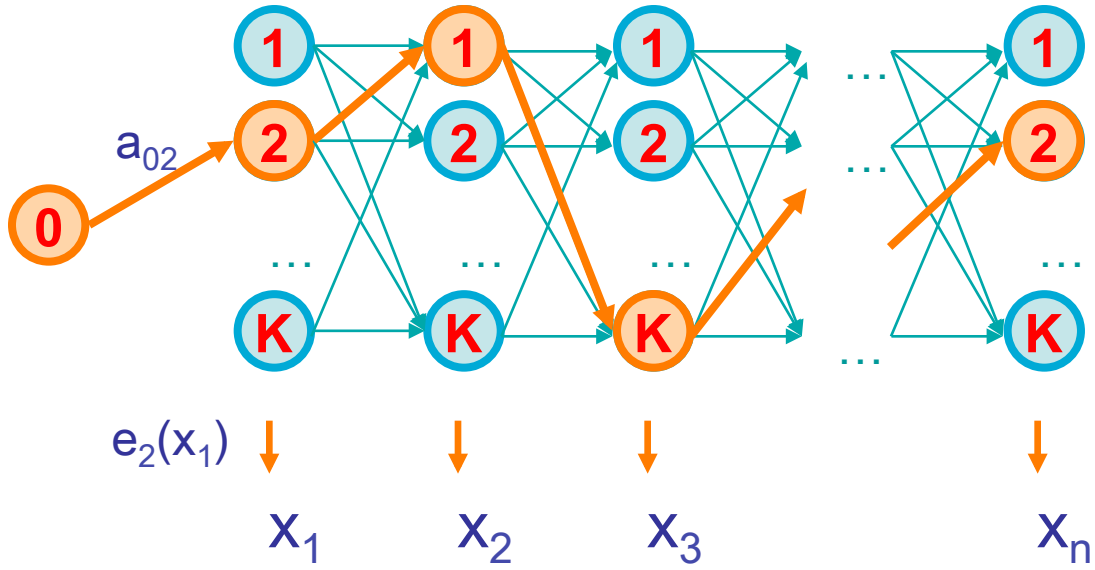


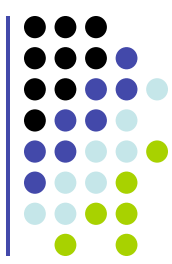


Generating a sequence by the model

Given a HMM, we can generate a sequence of length n as follows:

1. Start at state π_1 according to prob $a_{0\pi_1}$
2. Emit letter x_1 according to prob $e_{\pi_1}(x_1)$
3. Go to state π_2 according to prob $a_{\pi_1\pi_2}$
4. ... until emitting x_n

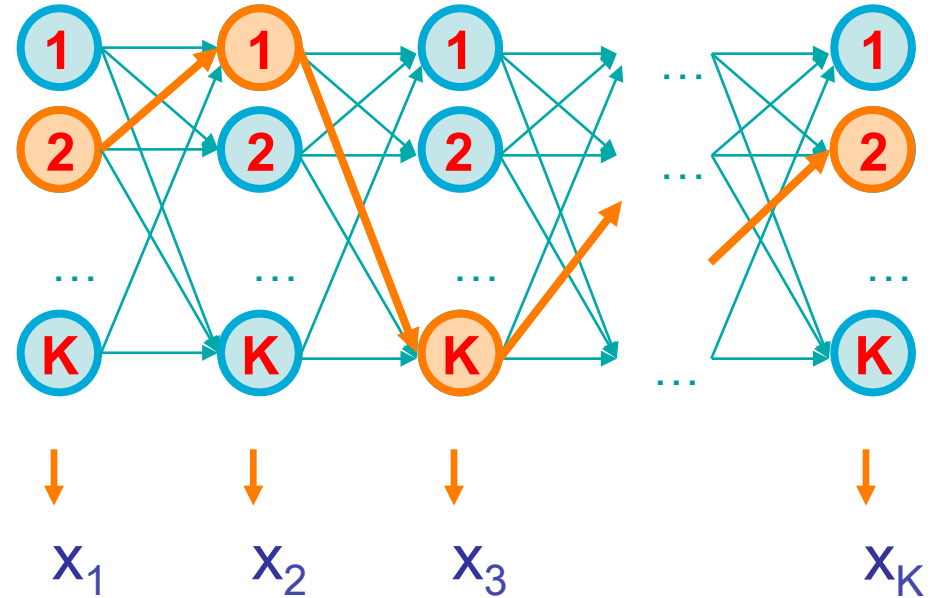




Likelihood of a parse

Given a sequence $\mathbf{x} = x_1 \dots x_N$
and a parse $\pi = \pi_1, \dots, \pi_N$,

To find how likely this scenario is:
(given our HMM)



$$P(\mathbf{x}, \pi) = P(x_1, \dots, x_N, \pi_1, \dots, \pi_N) =$$

$$P(x_N | \pi_N) P(\pi_N | \pi_{N-1}) \dots P(x_2 | \pi_2) P(\pi_2 | \pi_1) P(x_1 | \pi_1) P(\pi_1) =$$
$$a_{0\pi_1} a_{\pi_1\pi_2} \dots a_{\pi_{N-1}\pi_N} e_{\pi_1}(x_1) \dots e_{\pi_N}(x_N)$$



Likelihood of a parse

A compact way to write

$$a_{0\pi_1} a_{\pi_1\pi_2} \dots a_{\pi_{N-1}\pi_N} e_{\pi_1}(x_1) \dots e_{\pi_N}(x_N)$$

Enumerate all parameters a_{ij} and $e_i(b)$; n params

Example:

$$a_{0Fair} : \theta_1; a_{0Loaded} : \theta_2; \dots e_{Loaded}(6) = \theta_{18}$$

Then, count in x and π the # of times each parameter $j = 1, \dots, n$ occurs

$$F(j, x, \pi) = \# \text{ parameter } \theta_j \text{ occurs in } (x, \pi)$$

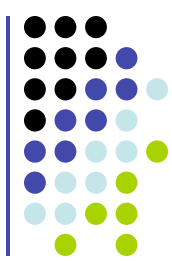
(call $F(.,.,.)$ the **feature counts**) Then,

$$P(x, \pi) = \prod_{j=1 \dots n} \theta_j^{F(j, x, \pi)} = \exp\left[\sum_{j=1 \dots n} \log(\theta_j) \times F(j, x, \pi)\right]$$

Given a sequence $x = x_1 \dots x_N$ and a parse $\pi = \pi_1, \dots, \pi_N$,

To find how likely this scenario (given our HMM)

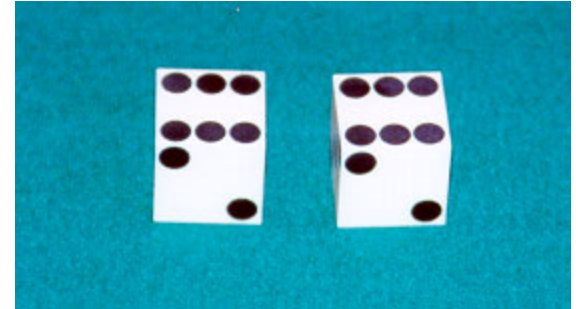
$$P(x, \pi) = P(x_1, \dots, x_N, \pi_1, \dots, \pi_N) = a_{0\pi_1} a_{\pi_1\pi_2} \dots a_{\pi_{N-1}\pi_N} e_{\pi_1}(x_1) \dots e_{\pi_N}(x_N)$$



Example: the dishonest casino

Let the sequence of rolls be:

$$x = 1, 2, 1, 5, 6, 2, 1, 5, 2, 4$$



Then, what is the likelihood of

$$\pi = \text{Fair, Fair, Fair, Fair, Fair, Fair, Fair, Fair, Fair, Fair?}$$

(say initial probs $a_{0\text{Fair}} = 1/2$, $a_{0\text{Loaded}} = 1/2$)

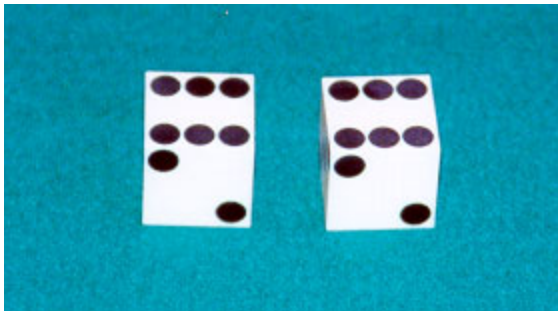
$$1/2 \times P(1 | \text{Fair}) P(\text{Fair} | \text{Fair}) P(2 | \text{Fair}) P(\text{Fair} | \text{Fair}) \dots P(4 | \text{Fair}) =$$

$$1/2 \times (1/6)^{10} \times (0.95)^9 = .00000000521158647211 \approx 0.5 \times 10^{-9}$$



Example: the dishonest casino

So, the likelihood the die is fair in this run is just 0.521×10^{-9}



What is the likelihood of

π = Loaded, Loaded, Loaded, Loaded, Loaded, Loaded, Loaded, Loaded, Loaded, Loaded, Loaded?

$$\frac{1}{2} \times P(1 \mid \text{Loaded}) P(\text{Loaded, Loaded}) \dots P(4 \mid \text{Loaded}) =$$

$$\frac{1}{2} \times (1/10)^9 \times (1/2)^1 (0.95)^9 = .00000000015756235243 \approx 0.16 \times 10^{-9}$$

Therefore, it is somewhat more likely that all the rolls are done with the fair die, than that they are all done with the loaded die



Example: the dishonest casino

Let the sequence of rolls be:

$$x = 1, 6, 6, 5, 6, 2, 6, 6, 3, 6$$

Now, what is the likelihood $\pi = F, F, \dots, F$?

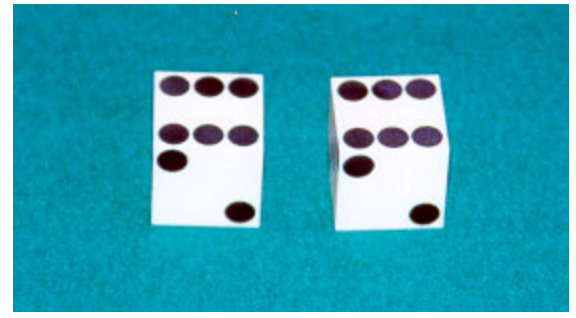
$$\frac{1}{2} \times (1/6)^{10} \times (0.95)^9 \approx 0.5 \times 10^{-9}, \text{ same as before}$$

What is the likelihood

$$\pi = L, L, \dots, L?$$

$$\frac{1}{2} \times (1/10)^4 \times (1/2)^6 (0.95)^9 = .00000049238235134735 \approx 0.5 \times 10^{-7}$$

So, it is 100 times more likely the die is loaded





Question # 1 – Evaluation

GIVEN

A sequence of rolls by the casino player

1245526462146146136136661664661636616366163616515615115146123562344

Prob = 1.3×10^{-35}

QUESTION

How likely is this sequence, given our model of how the casino works?

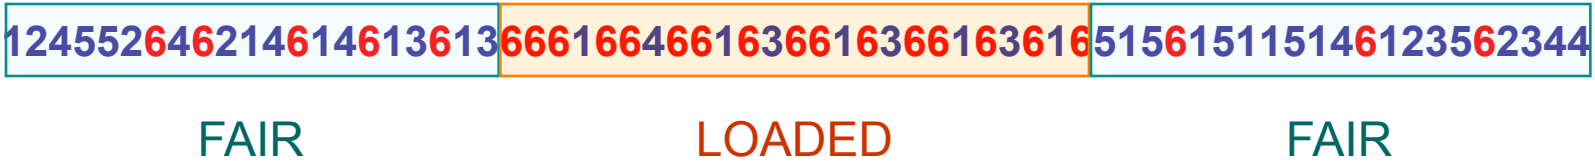
This is the **EVALUATION** problem in HMMs



Question # 2 – Decoding

GIVEN

A sequence of rolls by the casino player



QUESTION

What portion of the sequence was generated with the fair die, and what portion with the loaded die?

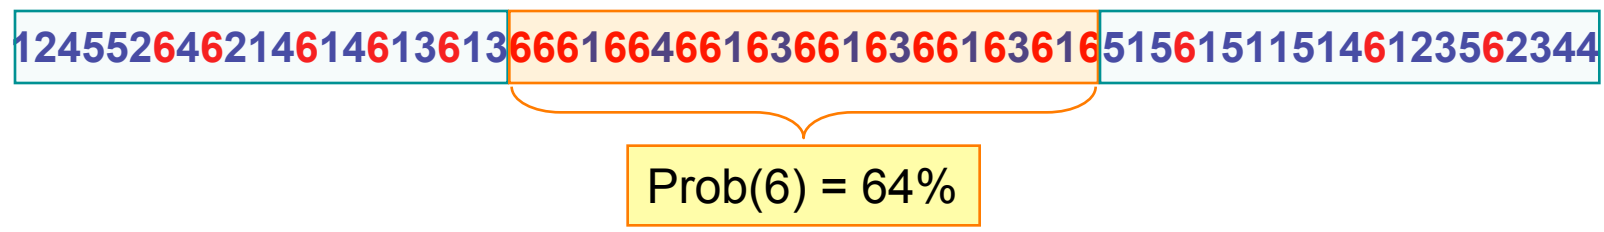
This is the **DECODING** question in HMMs



Question # 3 – Learning

GIVEN

A sequence of rolls by the casino player



QUESTION

How “loaded” is the loaded die? How “fair” is the fair die? How often does the casino player change from fair to loaded, and back?

This is the **LEARNING** question in HMMs



The three main questions on HMMs

1. Evaluation

GIVEN a HMM M , and a sequence x ,
FIND $\text{Prob}[x | M]$

2. Decoding

GIVEN a HMM M , and a sequence x ,
FIND the sequence π of states that maximizes $P[x, \pi | M]$

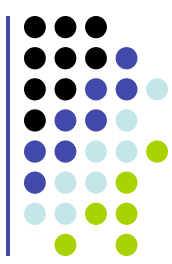
3. Learning

GIVEN a HMM M , with unspecified transition/emission probs.,
and a sequence x ,
FIND parameters $\theta = (e_i(\cdot), a_{ij})$ that maximize $P[x | \theta]$



Problem 1: Decoding

Find the most likely parse of a sequence



Decoding

GIVEN $x = x_1 x_2 \dots x_N$

Find $\pi = \pi_1, \dots, \pi_N$,
to maximize $P[x, \pi]$

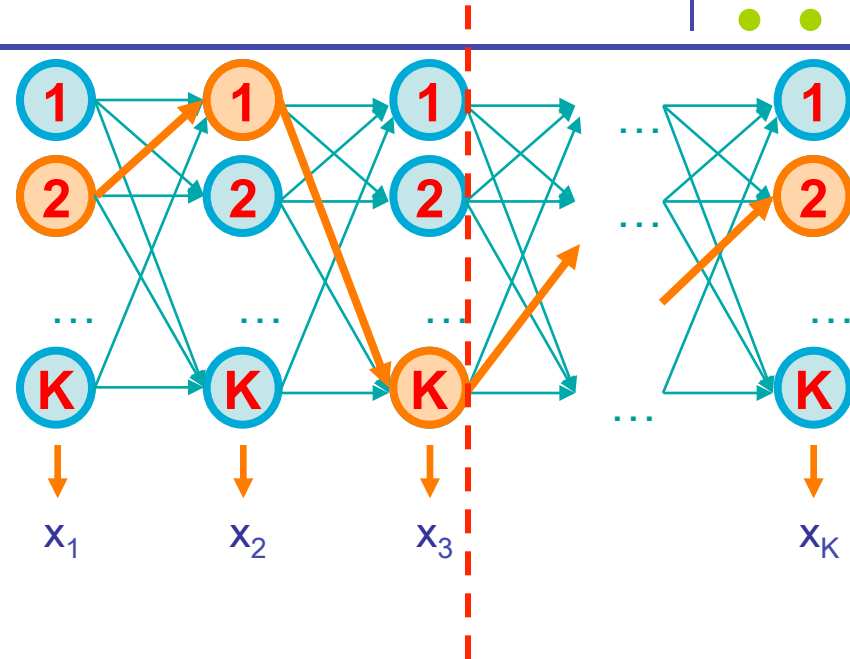
$\pi^* = \operatorname{argmax}_{\pi} P[x, \pi]$

Maximizes $a_{0\pi_1} e_{\pi_1}(x_1) a_{\pi_1\pi_2} \dots a_{\pi_{N-1}\pi_N} e_{\pi_N}(x_N)$

Dynamic Programming!

$V_k(i) = \max_{\{\pi_1 \dots \pi_{i-1}\}} P[x_1 \dots x_{i-1}, \pi_1, \dots, \pi_{i-1}, x_i, \pi_i = k]$

= Prob. of most likely sequence of states ending at
state $\pi_i = k$



Given that we end up in
state k at step i ,
maximize product to the
left and right



Decoding – main idea

Inductive assumption: Given that for all states k ,
and for a fixed position i ,

$$V_k(i) = \max_{\{\pi_1 \dots \pi_{i-1}\}} P[x_1 \dots x_{i-1}, \pi_1, \dots, \pi_{i-1}, x_i, \pi_i = k]$$

What is $V_l(i+1)$?

From definition,

$$\begin{aligned} V_l(i+1) &= \max_{\{\pi_1 \dots \pi_i\}} P[x_1 \dots x_i, \pi_1, \dots, \pi_i, x_{i+1}, \pi_{i+1} = l] \\ &= \max_{\{\pi_1 \dots \pi_i\}} P(x_{i+1}, \pi_{i+1} = l \mid x_1 \dots x_i, \pi_1, \dots, \pi_i) P[x_1 \dots x_i, \pi_1, \dots, \pi_i] \\ &= \max_{\{\pi_1 \dots \pi_i\}} P(x_{i+1}, \pi_{i+1} = l \mid \pi_i) P[x_1 \dots x_{i-1}, \pi_1, \dots, \pi_{i-1}, x_i, \pi_i] \\ &= \max_k [P(x_{i+1}, \pi_{i+1} = l \mid \pi_i = k) \max_{\{\pi_1 \dots \pi_{i-1}\}} P[x_1 \dots x_{i-1}, \pi_1, \dots, \pi_{i-1}, x_i, \pi_i = k]] \\ &= \max_k [P(x_{i+1} \mid \pi_{i+1} = l) P(\pi_{i+1} = l \mid \pi_i = k) V_k(i)] \\ &= e_l(x_{i+1}) \max_k a_{kl} V_k(i) \end{aligned}$$



The Viterbi Algorithm

Input: $x = x_1 \dots x_N$

Initialization:

$$V_0(0) = 1$$

(0 is the imaginary first position)

$$V_k(0) = 0, \text{ for all } k > 0$$

Iteration:

$$V_j(i) = e_j(x_i) \times \max_k a_{kj} V_k(i-1)$$

$$\text{Ptr}_j(i) = \text{argmax}_k a_{kj} V_k(i-1)$$

Termination:

$$P(x, \pi^*) = \max_k V_k(N)$$

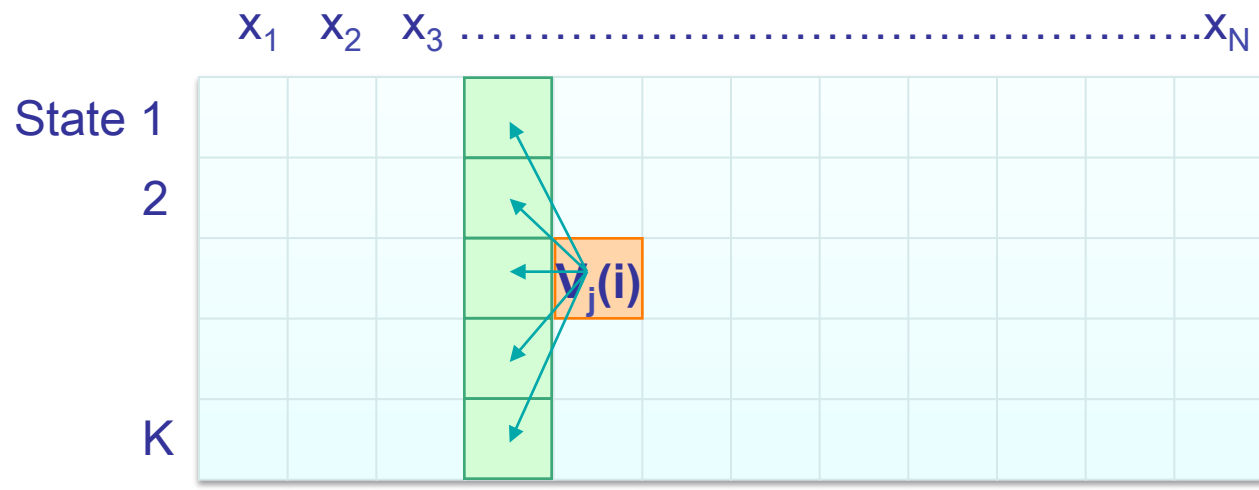
Traceback:

$$\pi_N^* = \text{argmax}_k V_k(N)$$

$$\pi_{i-1}^* = \text{Ptr}_{\pi_i^*}(i)$$



The Viterbi Algorithm



Similar to “aligning” a set of states to a sequence

Time:

$$O(K^2N)$$

Space:

$$O(KN)$$



Viterbi Algorithm – a practical detail

Underflows are a significant problem

$$P[\mathbf{x}_1, \dots, \mathbf{x}_i, \pi_1, \dots, \pi_i] = a_{0\pi_1} a_{\pi_1\pi_2} \dots a_{\pi_i} e_{\pi_1}(\mathbf{x}_1) \dots e_{\pi_i}(\mathbf{x}_i)$$

These numbers become extremely small – underflow

Solution: Take the logs of all values

$$V_i(i) = \log e_k(\mathbf{x}_i) + \max_k [V_k(i-1) + \log a_{ki}]$$



Example

Let x be a long sequence with a portion of $\sim 1/6$ 6's,
followed by a portion of $\sim 1/2$ 6's...

$x = 123456123456\dots123456$ $6626364656\dots1626364656$

Then, it is not hard to show that optimal parse is (exercise):

$FFF\dots\dots F$ $LLL\dots\dots L$

6 characters "123456" parsed as F, contribute $.95^6 \times (1/6)^6 = 1.6 \times 10^{-5}$
parsed as L, contribute $.95^6 \times (1/2)^1 \times (1/10)^5 = 0.4 \times 10^{-5}$

"162636" parsed as F, contribute $.95^6 \times (1/6)^6 = 1.6 \times 10^{-5}$
parsed as L, contribute $.95^6 \times (1/2)^3 \times (1/10)^3 = 9.0 \times 10^{-5}$



Problem 2: Evaluation

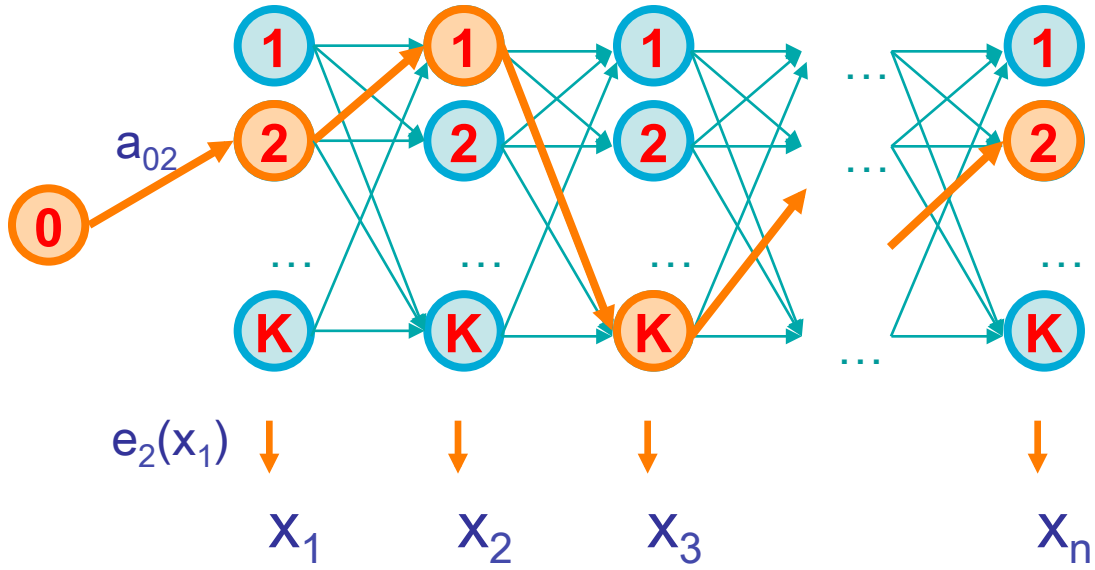
Find the likelihood a sequence
is generated by the model



Generating a sequence by the model

Given a HMM, we can generate a sequence of length n as follows:

1. Start at state π_1 according to prob $a_{0\pi_1}$
2. Emit letter x_1 according to prob $e_{\pi_1}(x_1)$
3. Go to state π_2 according to prob $a_{\pi_1\pi_2}$
4. ... until emitting x_n





A couple of questions

Given a sequence x,

- What is the probability that
- Given a position i, what is the

$$\begin{aligned}
 P(\text{box: FFFFFFFFFFFF}) &= \\
 &(1/6)^{11} * 0.95^{12} = \\
 &2.76^{-9} * 0.54 = \\
 &1.49^{-9} \\
 \\
 P(\text{box: LLLLLLLLLLLL}) &= \\
 &[(1/2)^6 * (1/10)^5] * 0.95^{10} * 0.05^2 = \\
 &1.56 * 10^{-7} * 1.5^{-3} = \\
 &0.23^{-9}
 \end{aligned}$$

Example: the dishonest ca

Say x = 12341...231 **62616364616** 234112...21341

F
F

Most likely path: $\pi = FF \dots F$

(too “unlikely” to transition $F \rightarrow L \rightarrow F$)

However: marked letters more likely to be L than unmarked letters



Evaluation

We will develop algorithms that allow us to compute:

$P(x)$ Probability of x given the model

$P(x_i \dots x_j)$ Probability of a substring of x given the model

$P(\pi_i = k \mid x)$ “**Posterior**” probability that the i^{th} state is k , given x

A more refined measure of which states x may be in



The Forward Algorithm

We want to calculate

$P(x)$ = probability of x , given the HMM

Sum over all possible ways of generating x :

$$P(x) = \sum_{\pi} P(x, \pi) = \sum_{\pi} P(x | \pi) P(\pi)$$

To avoid summing over an exponential number of paths π , define

$$f_k(i) = P(x_1 \dots x_i, \pi_i = k) \quad (\text{the } \textbf{forward} \text{ probability})$$

“generate i first characters of x and end up in state k ”



The Forward Algorithm – derivation

Define the forward probability:

$$\begin{aligned} f_k(i) &= P(x_1 \dots x_i, \pi_i = k) \\ &= \sum_{\pi_1 \dots \pi_{i-1}} P(x_1 \dots x_{i-1}, \pi_1, \dots, \pi_{i-1}, \pi_i = k) e_k(x_i) \\ &= \sum_l \sum_{\pi_1 \dots \pi_{i-2}} P(x_1 \dots x_{i-1}, \pi_1, \dots, \pi_{i-2}, \pi_{i-1} = l) a_{lk} e_k(x_i) \\ &= \sum_l P(x_1 \dots x_{i-1}, \pi_{i-1} = l) a_{lk} e_k(x_i) \\ &= e_k(x_i) \sum_l f_l(i-1) a_{lk} \end{aligned}$$



The Forward Algorithm

We can compute $f_k(i)$ for all k, i , using dynamic programming!

Initialization:

$$f_0(0) = 1$$

$$f_k(0) = 0, \text{ for all } k > 0$$

Iteration:

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

Termination:

$$P(x) = \sum_k f_k(N)$$



Relation between Forward and Viterbi

VITERBI

Initialization:

$$V_0(0) = 1$$
$$V_k(0) = 0, \text{ for all } k > 0$$

Iteration:

$$V_j(i) = e_j(x_i) \max_k V_k(i - 1) a_{kj}$$

Termination:

$$P(x, \pi^*) = \max_k V_k(N)$$

FORWARD

Initialization:

$$f_0(0) = 1$$
$$f_k(0) = 0, \text{ for all } k > 0$$

Iteration:

$$f_l(i) = e_l(x_i) \sum_k f_k(i - 1) a_{kl}$$

Termination:

$$P(x) = \sum_k f_k(N)$$



Motivation for the Backward Algorithm

We want to compute

$$P(\pi_i = k \mid x),$$

the probability distribution on the i^{th} position, given x

We start by computing

$$\begin{aligned} P(\pi_i = k, x) &= P(x_1 \dots x_i, \pi_i = k, x_{i+1} \dots x_N) \\ &= P(x_1 \dots x_i, \pi_i = k) P(x_{i+1} \dots x_N \mid x_1 \dots x_i, \pi_i = k) \\ &= \boxed{P(x_1 \dots x_i, \pi_i = k)} \boxed{P(x_{i+1} \dots x_N \mid \pi_i = k)} \end{aligned}$$

Forward, $f_k(i)$ **Backward, $b_k(i)$**

Then, $P(\pi_i = k \mid x) = P(\pi_i = k, x) / P(x)$



The Backward Algorithm – derivation

Define the backward probability:

$$\begin{aligned} b_k(i) &= P(x_{i+1} \dots x_N \mid \pi_i = k) && \text{“starting from } i^{\text{th}} \text{ state = } k, \text{ generate rest of } x\text{”} \\ &= \sum_{\pi_{i+1} \dots \pi_N} P(x_{i+1}, x_{i+2}, \dots, x_N, \pi_{i+1}, \dots, \pi_N \mid \pi_i = k) \\ &= \sum_l \sum_{\pi_{i+1} \dots \pi_N} P(x_{i+1}, x_{i+2}, \dots, x_N, \pi_{i+1} = l, \pi_{i+2}, \dots, \pi_N \mid \pi_i = k) \\ &= \sum_l e_l(x_{i+1}) a_{kl} \sum_{\pi_{i+1} \dots \pi_N} P(x_{i+2}, \dots, x_N, \pi_{i+2}, \dots, \pi_N \mid \pi_{i+1} = l) \\ &= \sum_l e_l(x_{i+1}) a_{kl} \mathbf{b}_l(i+1) \end{aligned}$$



The Backward Algorithm

We can compute $b_k(i)$ for all k, i , using dynamic programming

Initialization:

$$b_k(N) = 1, \text{ for all } k$$

Iteration:

$$b_k(i) = \sum_l e_l(x_{i+1}) a_{kl} b_l(i+1)$$

Termination:

$$P(x) = \sum_l a_{0l} e_l(x_1) b_l(1)$$



Computational Complexity

What is the running time, and space required, for Forward, and Backward?

Time: $O(K^2N)$
Space: $O(KN)$

Useful implementation technique to avoid underflows

Viterbi: sum of logs
Forward/Backward: rescaling at each few positions by multiplying by a constant



Posterior Decoding

We can now calculate

$$P(\pi_i = k | x) = \frac{f_k(i) b_k(i)}{P(x)}$$

Then, we can ask

$$\begin{aligned} P(\pi_i = k | x) &= \\ P(\pi_i = k, x) / P(x) &= \\ P(x_1, \dots, x_i, \pi_i = k, x_{i+1}, \dots, x_n) / P(x) &= \\ P(x_1, \dots, x_i, \pi_i = k) P(x_{i+1}, \dots, x_n | \pi_i = k) / P(x) &= \\ f_k(i) b_k(i) / P(x) & \end{aligned}$$

What is the most likely state at position i of sequence x:

Define π^{\wedge} by Posterior Decoding:

$$\pi^{\wedge}_i = \operatorname{argmax}_k P(\pi_i = k | x)$$

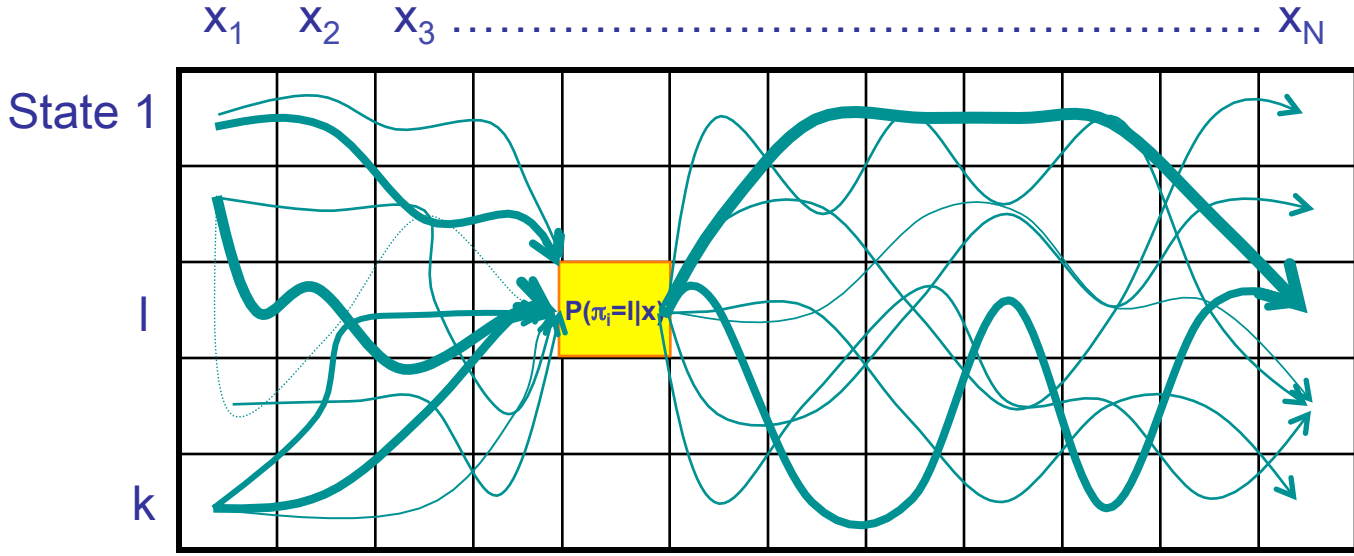


Posterior Decoding

- For each state,
 - Posterior Decoding gives us a curve of likelihood of state for each position
 - That is sometimes more informative than Viterbi path π^*
- Posterior Decoding may give an invalid sequence of states (of prob 0)
 - Why?



Posterior Decoding



- $$P(\pi_i = k | x) = \sum_{\pi} P(\pi | x) \mathbf{1}(\pi_i = k)$$
$$= \sum_{\{\pi: \pi[i] = k\}} P(\pi | x)$$

$\mathbf{1}(\psi) = 1$, if ψ is true
 0 , otherwise



Viterbi, Forward, Backward

VITERBI

Initialization:

$$V_0(0) = 1$$

$$V_k(0) = 0, \text{ for all } k > 0$$

Iteration:

$$V_i(i) = e_i(x_i) \max_k V_k(i-1) a_{ki}$$

Termination:

$$P(x, \pi^*) = \max_k V_k(N)$$

FORWARD

Initialization:

$$f_0(0) = 1$$

$$f_k(0) = 0, \text{ for all } k > 0$$

Iteration:

$$f_i(i) = e_i(x_i) \sum_k f_k(i-1) a_{ki}$$

Termination:

$$P(x) = \sum_k f_k(N)$$

BACKWARD

Initialization:

$$b_k(N) = 1, \text{ for all } k$$

Iteration:

$$b_i(i) = \sum_k e_i(x_{i+1}) a_{ki} b_k(i+1)$$

Termination:

$$P(x) = \sum_k a_{0k} e_k(x_1) b_k(1)$$