

# Algorithms in Bioinformatics

Randomized  
Algorithms  
Gibbs sampling  
& motif search

Recap

- Final implementation  
project

# Today: Randomization (Ch. 12)

We've used some probability,  
but haven't yet really  
focused on randomized  
algorithms.

Idea: Use randomness to:

- "foil" an adversary:  
choosing something randomly  
makes pathological cases  
less likely
- randomly sample:  
Markov-chain:  
sample huge space, but  
only small part of it
- hashing / finger-printing:
- load balancing

Note:

# Broad Categories

- ① • Las Vegas algorithms:
- always return correct answers
  - but varying run time

- ② • Monte Carlo algorithms:
- may produce incorrect or approximate solutions
  - very extensively used in bioinformatics though



# First look: Randomized quicksort.

Algorithm: input:  $S$

Choose a random element  $e$   
("pivot")

Determine

$$S_L = \{\text{elements of } S < e\}$$

$$S_G = \{\text{elements of } S \geq e\}$$

Recursively sort  $S_L + S_G$

Return  $\{\text{sorted } S_L\} ++ e$   
 $++ \{\text{sorted } S_G\}$

How to do runtime?

Let  $S_{(i)}$  be element of  
rank  $i$  in  $S$

(so  $S_{(1)}$  is smallest +  
 $S_{(n)}$  is largest)

Let  $X_{ij} = \begin{cases} 1 & \text{if } S_{(i)} + S_{(j)} \\ & \text{are compared} \\ 0 & \text{if not} \end{cases}$

indicator variable

Then:

$$\begin{aligned} \text{Total \# of comparisons} \\ = \sum_{i=1}^n \sum_{j>i} X_{ij} \end{aligned}$$

Our goal: expected # of comparisons

$$E\left[\sum_{i=1}^n \sum_{j>i} X_{ij}\right]$$

~~\*~~ → use linearity of expectation:


$$= \sum_{i=1}^n \sum_{j>i} E[X_{ij}]$$

If we let  $p_{ij}$  = prob.  $i$  &  $j$  are compared, then

$$\begin{aligned} E[X_{ij}] &= 1 \cdot p_{ij} + 0(1-p_{ij}) \\ &= p_{ij} \end{aligned}$$

Shifting our view:

View execution as a binary tree, where each node gets labeled with its pivot choice



Note: root value  $s$  is compared to everything

But: **nothing in  $S_L$  is ever compared to anything in  $S_G$**

## 2 observations:

as pivot

- $S_{(i)}$  &  $S_{(j)}$  are compared only if either is chosen before.  $\downarrow$   
any with  $S_{(l)}$   $i < l < j$ .


Why?



- Any of  $S_{(i)}, S_{(i+1)}, \dots, S_{(j)}$  is equally likely to be chosen as a first partition from this range.

So prob. that  $S_{(i)}$  or  $S_{(j)}$  is first is  $\equiv$

$$\frac{2}{j-i+1}$$



$$\text{So: } P_{ij} = \frac{2}{j-i+1}$$

$$E[\# \text{ comparisons}]$$

$$= \sum_{i=1}^n \sum_{\substack{j>i \\ j-i+1}}^n \cancel{P_{ij}} \frac{2}{j-i+1}$$

$$= 2 \sum_{i=1}^n \sum_{k=2}^{n-i+1} \frac{1}{k}$$

$$\leq 2 \sum_{i=1}^n \left( \sum_{k=1}^n \frac{1}{k} \right)$$

Def: The  $n$ th harmonic number,  $H_n$ , is defined as  $\sum_{k=1}^n \frac{1}{k}$

This is  $\approx \ln n + O(1)$

$\Rightarrow$  Quick sort runs in  $O(n \log n)$  expected time.

Key thing:

- This expected running time holds for every input.

Randomness depends upon the algorithm making random choices.

- Exact runtime will still vary - we are treating runtime as the random variable.

- This one is still always correct, though.  
(not an approximation)

(Las Vegas)

# Gibbs Sampling:

An older method, based on Markov chain Monte Carlo methods - 1953-ish.

First applied to motifs in 1993 by Lawrence et al.

## Recall: Motif Finding

Find an  $l$ -mer from each of  $t$  input sequences s.t. similarity is maximized

find a better alignment?

```

CGGGGCTATcCagCTGGGTCGTCACATCCCCTT...
TTTGAGGcTGCCCAATAAggGCAACTCCAAGCGGACAAA
GGATGgACTGATGCGGTTTGACGACCTA...
AAGGAaGCAAGCCAGGAGCCGCTTTGCTGG...
AATTTTCTAAAAAGATATAATGTGGTCCCTGGAACTTC
CTGCTGTACAACCTGAGATCATGCTGCATGCCaLTTTCAAC
TACATGATCTTTTGATGgcACTTTGGATGAGGGAATGATGC
    
```

We'll view it this way.

(a) Superposition of the seven highlighted 8-mers from figure 4.2 (d).

	A	T	C	A	G	C	T
Alignment	G	G	G	C	A	A	C
	A	T	G	G	A	A	T
	A	A	G	C	A	A	C
	T	T	G	G	A	A	C
	A	T	G	C	C	A	T
	A	T	G	G	C	A	C
	A	T	G	G	C	A	C
Profile	A	5	1	0	0	5	5
	T	1	5	0	0	0	1
	G	1	1	6	3	0	1
	C	0	0	1	4	2	0
Consensus	A	T	G	C	A	A	C

A	.72	.14	0	0	.72	.72	0	0
T	.14	.72	0	0	0	.14	.14	.86
G	.14	.14	.86	.44	0	.14	0	0
C	0	0	.14	.56	.28	0	.86	.14

(b) The alignment matrix, profile matrix and consensus string formed from the 8-mers starting at positions  $s = (8, 19, 3, 5, 31, 27, 15)$  in figure 4.2 (d).

ch 11

Figure 4.3 From DNA sample, to alignment matrix, to profile, and finally, to consensus string. If  $s = (8, 19, 3, 5, 31, 27, 15)$  is an array of starting positions for 8-mers in figure 4.2 (d), then  $Score(s) = 5 + 5 + 6 + 4 + 5 + 5 + 6 + 6 = 42$ .

Given a profile P:

A	.72	.14	0	0	.72	.72	0	0
T	.14	.72	0	0	0	.14	.14	.86
G	.14	.14	.86	.44	0	.14	0	0
C	0	0	.14	.56	.28	0	.86	.14

• arbitrary l-mer:  $a_1 \dots a_l$ ,  
let  $P(a|P) = \prod_{i=1}^l p_{a_i, i}$

This is the probability that a was generated by P.

Example:  $a = \text{ATGCAACT}$   
consensus string

$$P(a|P) = 0.72 \times 0.72 \times 0.14 \dots$$

$$\approx 9.6 \times 10^{-2}$$

$$P(\text{TACGCGTC}|P)$$

=

$$\approx 9.3 \times 10^{-7}$$



So: you can evaluate the probability of each  $l$ -mer and find the most likely one.

- called the P-most <sup>we don't know this</sup> probable  $l$ -mer

Motivates a random approach:

- Start with random starting positions
- Try to greedily improve

```
GREEDYPROFILEMOTIFSEARCH( $DNA, t, n, l$ )
```

- 1 Randomly select starting positions  $s = (s_1, \dots, s_t)$  in  $DNA$
- 2 Form profile  $P$  from  $s$
- 3  $bestScore \leftarrow 0$
- 4 **while**  $Score(s, DNA) > bestScore$
- 5      $bestScore \leftarrow Score(s, DNA)$
- 6     **for**  $i \leftarrow 1$  to  $t$
- 7         Find a P-most probable  $l$ -mer  $a$  from the  $i$ th sequence
- 8          $s_i \leftarrow$  starting position of  $a$
- 9 **return**  $bestScore$

Problem: It jumps around in large search space.

The last algorithm changes the starting positions in each iteration

Gibbs sampling moves more slowly:

In each iteration, discard one  $l$ -mer & replace with a new one.


1. Randomly select starting positions  $s = (s_1, \dots, s_t)$  in DNA and form the set of  $l$ -tuples starting at these positions.
- ~~2.~~ 2. Randomly choose one sequence out of  $t$  DNA sequences.
3. Create a profile  $\mathbf{P}$  from the  $l$ -mers in the remaining  $t - 1$  sequences.
4. For each position  $i$  in the chosen DNA sequence, calculate the probability  $p_i$  that the  $l$ -mer starting at this position is generated by profile  $\mathbf{P}$  ( $1 \leq i \leq n - l + 1$ ).
5. Choose the new starting position in the chosen DNA sequence randomly, according to the distribution proportional to  $(p_1, p_2, \dots, p_{n-l+1})$ .
6. Repeat until convergence.<sup>4</sup>

↑  
we don't  
address this


# Illustration (from MIT demo):

## Gibbs Sampling Algorithm I


1. Select a **random** position in each sequence




2. Build a weight matrix




3. Select a sequence at random




#### 4. Score possible sites in the sequence using weight matrix



#### 5. Sample a new site proportional to likelihood and update motif instances



#### 6. Update weight matrix




Convergence:  $\Delta_{sites} = 0$  or  $\Delta\theta \approx 0$

The hope:

By moving around (& running several times), we can find a good maximum.

Problem (as always):

What the motif landscape might look like



So - this is a Monte Carlo.

A particular weakness:

If the nucleotide distribution is skewed, i.e.

- A has 70% frequency

then search may lead to patterns composed of most frequent, which may not be biologically relevant.

Some use relative entropy to address this:

$$\text{entropy} = \sum_{j=1}^l \sum_{r \in \{A, T, C, G\}} P_{rj} \log \frac{P_{rj}}{b_r}$$

where  $P_{rj}$  = frequency of nucleotide  $r$  in position  $j$  of the alignment

# Random Projections

If  $l$ -mer is in each DNA strand with mutations, one idea is to try to focus on un-mutated spots.

Gaps may be in different spots, though:

CG ~~X~~ CAT ~~X~~ G

CG ~~X~~ CA ~~X~~ AG

~~C~~ ~~X~~ TCA ~~X~~ AG

↳ "Consensus" gapped pattern is

C ~~X~~ ~~X~~ CA ~~X~~ ~~X~~ G

However, these 4 spots are not known!

Random projection: pick them randomly

$(k, l)$ -template  $t$ :  
any  $k$  integers  $1 \leq t_1 < \dots < t_k \leq l$

For an  $l$ -mer  $a_1, \dots, a_l$ ,

Projection( $a, t$ )

= concatenation of  
nucleotides from the  
template

Example:  $a = \text{ATGCATI}$

$t = (2, 5, 7)$

Proj( $a, t$ ) = **TAT**

Idea: choose a random  $t$   
project every  $l$ -mer with it  
record via hash table

Expect likely motif will  
lead to higher counts.



RANDOMPROJECTIONS( $DNA, t, n, l, k, \theta, m$ )

```
1 create a  $t \times n$  array motifs and fill it with zeros
2 for  $m$  iterations
3 create a table Bins of size  $4^k$  and fill it with zeros
4  $r \leftarrow$  a random  $(k, l)$ -template.
5 for  $i \leftarrow 1$  to  $t$ 
6     for  $j \leftarrow 1$  to  $n - l + 1$ 
7          $a \leftarrow$   $j$ th  $l$ -mer in  $i$ th  $DNA$  sequence
8         Bins( $Projection(a, r)$ ) = Bins( $Projection(a, r)$ ) + 1
9     for  $i \leftarrow 1$  to  $t$ 
10        for  $j \leftarrow 1$  to  $n - l + 1$ 
11             $a \leftarrow$   $j$ th  $l$ -mer in  $i$ th  $DNA$  sequence
12            if Bins( $Projection(a, r)$ ) >  $\theta$ 
13                 $motifs_{i,j} \leftarrow motifs_{i,j} + 1$ 
14 for  $i \leftarrow 1$  to  $t$ 
15      $s_i \leftarrow$  Index of the largest element in row  $i$  of motifs.
16 return  $s$ 
```

## Notes:

- Parameters:

Selects  $m$  random  
 $(k, l)$ -templates

Aggregates data for all  
 $\cup \cup m$  of them

-  $\theta$  is significance threshold

- Also need table: **Bins**( $x$ )  
counts # of  $l$ -mers  
s.t.  $Proj(a, r) = x$

No guarantee.

However, can show that  
it 'works' with  
high probability.

(given good parameters)

Practical version by  
Buhler + Tompa is  
a bit more complex,  
but uses a heuristic  
method to choose  
the final  $(s_1, \dots, s_t)$