

Welcome to CS262: Computational Genomics

Instructor:

Serafim Batzoglou

TA:


Paul Chen

email: cs262-win2015-staff@lists.stanford.edu

Tuesdays & Thursdays 12:50-2:05pm
Clark S361

<http://cs262.stanford.edu>





Goals of this course

- Introduction to Computational Biology & Genomics
 - Basic concepts and scientific questions
 - Why does it matter?
 - Basic biology for computer scientists
 - In-depth coverage of algorithmic techniques
 - Current active areas of research
- Useful algorithms
 - Dynamic programming
 - String algorithms
 - HMMs and other graphical models for sequence analysis




Topics in CS262

Part 1: Basic Algorithms

- Dynamic Programming & sequence alignment
- HMMs, CRFs & sequence modeling
- Sequence indexing; Burrows-Wheeler transform, De Bruijn graphs


Part 2: Topics in computational genomics and areas of active research

- DNA sequencing and assembly
- Comparative genomics
- Human genome resequencing
 - Alignment
 - Compression
 - Human genome variation
- Cancer genomics
- Functional genomics
- Population genomics



Course responsibilities

- Homeworks
 - 4 challenging problem sets, 4-5 problems/pset
 - Due at beginning of class
 - Up to 3 late days (24-hr periods) for the quarter
 - Collaboration allowed – please give credit
 - Teams of 2 or 3 students
 - Individual writeups
 - If individual (no team) then drop score of worst problem per problem set
- (Optional) Scribing
 - Due one week after the lecture, except special permission
 - Scribing grade replaces 2 lowest problems from all problem sets
 - First-come first-serve, email staff list to sign up





Reading material


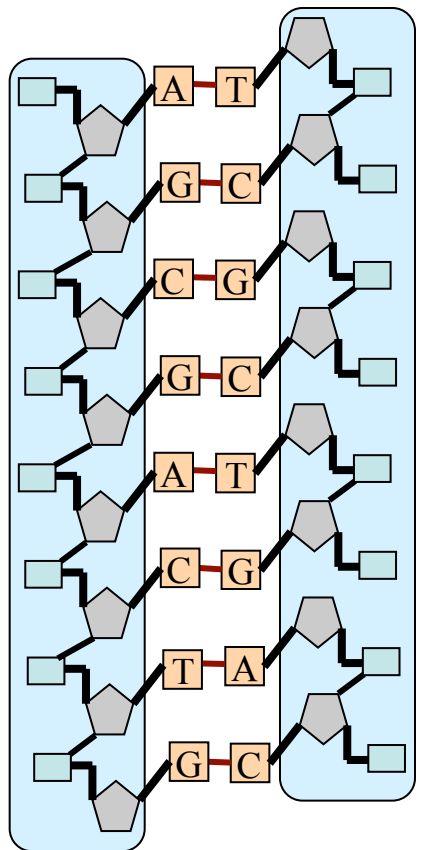
- Main Reading:
 - Lecture notes
 - Papers
- Optional:
 - “**Biological sequence analysis**” by Durbin, Eddy, Krogh, Mitchison
 - Chapters 1-4, 6, 7-8, 9-10



Birth of Molecular Biology



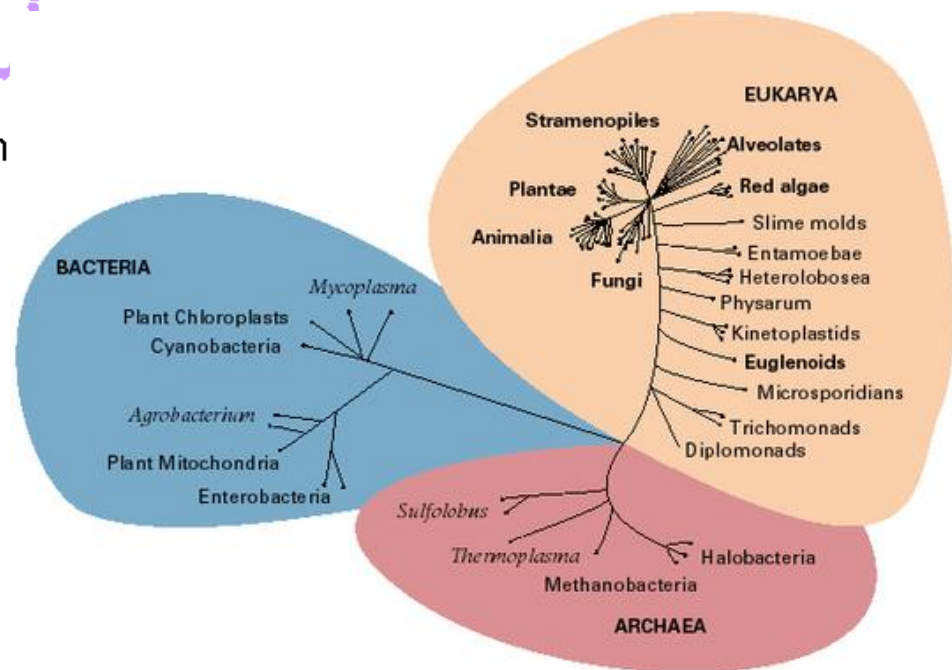

DNA




Courtesy of Cold Spring Harbor Laboratory Archives. Noncommercial, educational use only.




Genetics in the 20th Century





Human Genome Project



3 billion basepairs

\$3 billion

1990: Start


2000: Bill Clinton:

“most important scientific discovery in the 20th century”

2001: Draft

2003: Finished

now what?




Sequencing Growth

Cost of one human genome

- 2004: \$30,000,000
- 2008: \$100,000
- 2010: \$10,000
- **2014: "\$1,000"**
- ???: \$300



How much would you pay for a smartphone?





Uses of Genomes


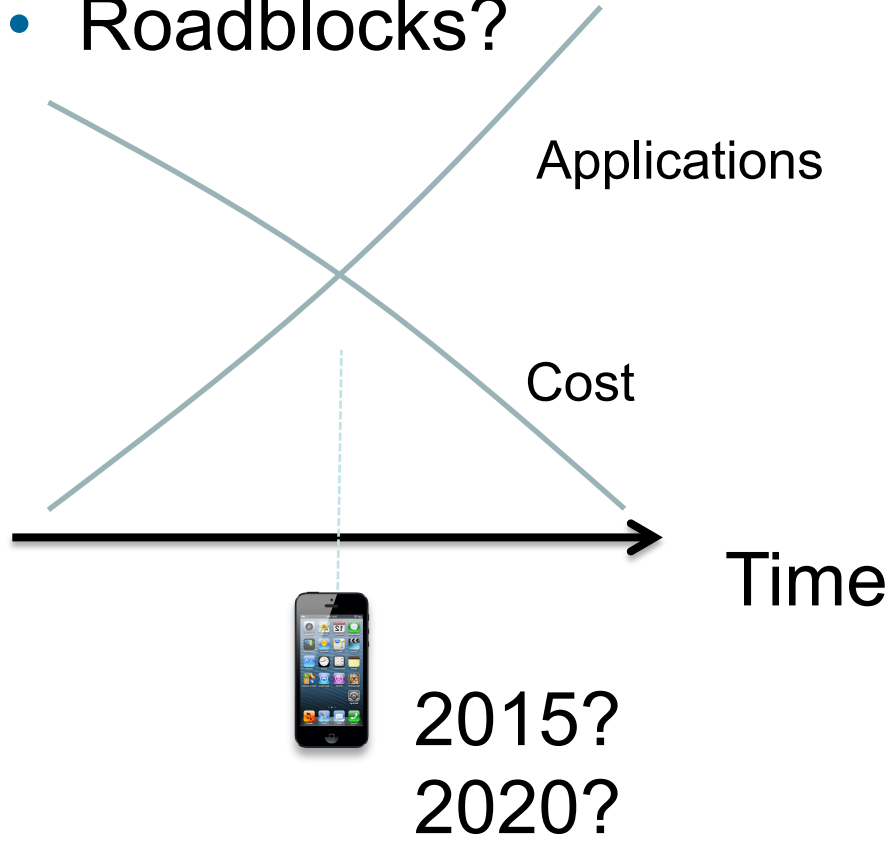
- Medicine
 - Prenatal/Mendelian diseases
 - Drug dosage (eg. Warfarin)
 - Disease risk
 - Diagnosis of infections
 - ...
- Ancestry
- Genealogy
- Nutrition?
- Psychology?
- Baby Engineering???
- **Ethical Issues**





How soon will we all be sequenced?

- Cost
- Killer apps
- Roadblocks?



Intro to Biology






Sequence Alignment






Evolution






Evolution at the DNA level




SEQUENCE EDITS

REARRANGEMENTS






Evolutionary Rates



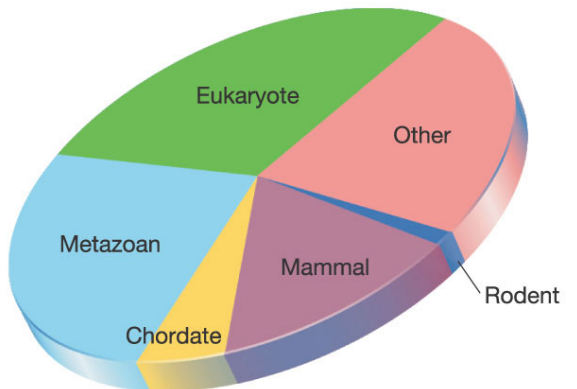


Sequence conservation implies function



Alignment is the key to

- Finding important regions
- Determining function
- Uncovering evolutionary events





Sequence Alignment

AGGCTATCACCTGACCTCCAGGCCGATGCCC
TAGCTATCACGACCGCGGGTCGATTTGCCCGAC

-AGGCTATCACCTGACCTCCAGGCCGA--TGCCC---
TAG-CTATCAC--GACCGC--GGTCGATTTGCCCGAC

Definition

Given two strings $x = x_1x_2 \dots x_M$, $y = y_1y_2 \dots y_N$,

an alignment is an assignment of gaps to positions $0, \dots, N$ in x , and $0, \dots, N$ in y , so as to line up each letter in one sequence with either a letter, or a gap in the other sequence



What is a good alignment?

AGGCTAGTT ,
AGCGAAGTTT

AGGCTAGTT-
AGCGAAGTTT

6 matches, 3 mismatches, 1 gap

AGGCTA-GTT-
AG-CGAAGTTT

7 matches, 1 mismatch, 3 gaps

AGGC-TA-GTT-
AG-CG-AAGTTT

7 matches, 0 mismatches, 5 gaps



Scoring Function

- Sequence edits:

- Mutations
- Insertions
- Deletions

AGGCCTC

AGGACTC

AGGGCCTC

AGG . CTC

Alternative definition:
minimal edit distance

“Given two strings x, y , find minimum # of edits (insertions, deletions, mutations) to transform one string to the other”

Scoring Function:

Match: +m
Mismatch: -s
Gap: -d

$$\text{Score } F = (\# \text{ matches}) \times m - (\# \text{ mismatches}) \times s - (\# \text{ gaps}) \times d$$



How do we compute the best alignment?

AGTGCCCTGGAACCCTGACGGTGGGTCACAAAACCTTCTGGA



Too many possible alignments:

$$\gg 2^N$$

(exercise)



Alignment is additive

Observation:

The score of aligning

$x_1 \dots x_M$

$y_1 \dots y_N$

is additive

Say that
aligns to

$x_1 \dots x_i$


$x_{i+1} \dots x_M$

$y_1 \dots y_j$

$y_{j+1} \dots y_N$

The two scores add up:

$$F(x[1:M], y[1:N]) = F(x[1:i], y[1:j]) + F(x[i+1:M], y[j+1:N])$$



Dynamic Programming

- There are only a polynomial number of subproblems
 - Align $x_1 \dots x_i$ to $y_1 \dots y_j$
- Original problem is one of the subproblems
 - Align $x_1 \dots x_M$ to $y_1 \dots y_N$
- Each subproblem is easily solved from smaller subproblems
 - We will show next
- Then, we can apply *Dynamic Programming!!!*

Let

$F(i, j)$ = optimal score of aligning
 $x_1 \dots x_i$
 $y_1 \dots y_j$

F is the DP “Matrix” or “Table”

“Memoization”



Dynamic Programming (cont'd)

Notice three possible cases:

1. x_i aligns to y_j

$$\begin{array}{l} x_1 \dots x_{i-1} \quad x_i \\ y_1 \dots y_{j-1} \quad y_j \end{array}$$
$$F(i, j) = F(i - 1, j - 1) + \begin{cases} m, & \text{if } x_i = y_j \\ -s, & \text{if not} \end{cases}$$

2. x_i aligns to a gap

$$\begin{array}{l} x_1 \dots x_{i-1} \quad x_i \\ y_1 \dots y_j \quad - \end{array}$$
$$F(i, j) = F(i - 1, j) - d$$

3. y_j aligns to a gap

$$\begin{array}{l} x_1 \dots x_i \quad - \\ y_1 \dots y_{j-1} \quad y_j \end{array}$$
$$F(i, j) = F(i, j - 1) - d$$



Dynamic Programming (cont'd)

How do we know which case is correct?

Inductive assumption:

$F(i, j - 1)$, $F(i - 1, j)$, $F(i - 1, j - 1)$ are optimal

Then,

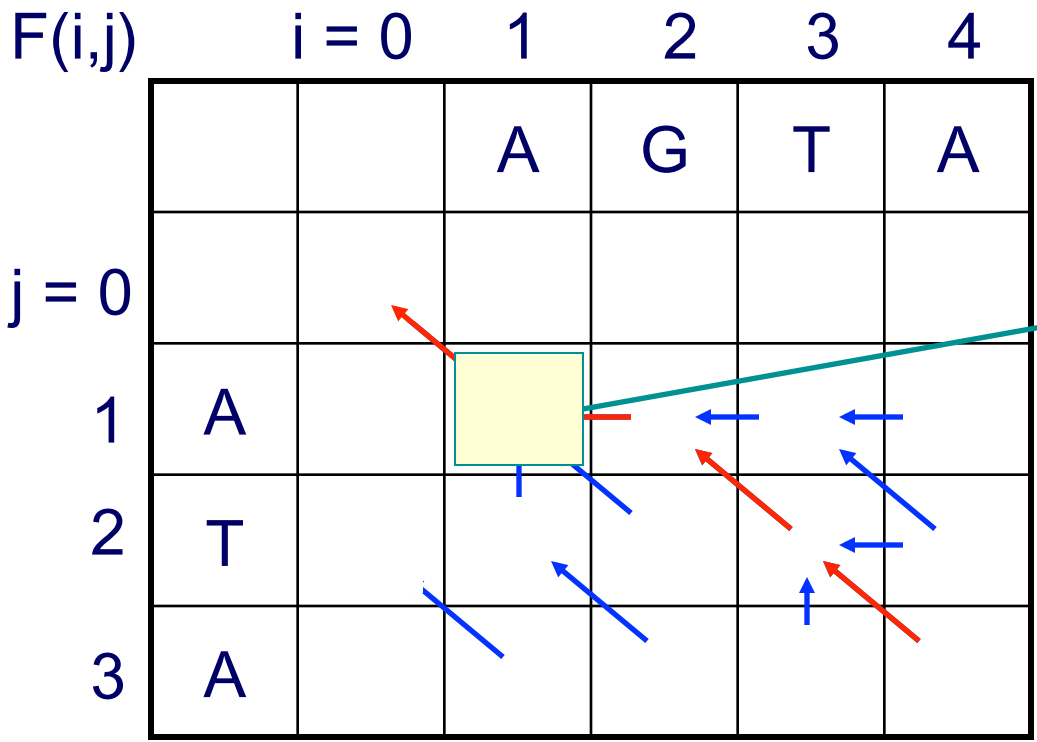
$$F(i, j) = \max \begin{cases} F(i - 1, j - 1) + s(x_i, y_j) \\ F(i - 1, j) - d \\ F(i, j - 1) - d \end{cases}$$

Where $s(x_i, y_j) = m$, if $x_i = y_j$; $-s$, if not



Example

$x = \text{AGTA}$
 $y = \text{ATA}$




Procedure to output Alignment

- Follow the backpointers
- When diagonal, OUTPUT x_i, y_j
- When up, OUTPUT y_j
- When left, OUTPUT x_i

AGTA
A - TA



The Needleman-Wunsch Matrix



Every nondecreasing path

from (0,0) to (M, N)

corresponds to an alignment of the two sequences

An optimal alignment is composed of optimal subalignments



The Needleman-Wunsch Algorithm

Initialization.

$$\begin{aligned} F(0, 0) &= 0 \\ F(0, j) &= -j \times d \\ F(i, 0) &= -i \times d \end{aligned}$$


Main Iteration. Filling-in partial alignments

For each $i = 1 \dots M$
For each $j = 1 \dots N$

$$F(i, j) = \max \begin{cases} F(i-1, j-1) + s(x_i, y_j) & \text{[case 1]} \\ F(i-1, j) - d & \text{[case 2]} \\ F(i, j-1) - d & \text{[case 3]} \end{cases}$$

$$Ptr(i, j) = \begin{cases} \text{DIAG,} & \text{if [case 1]} \\ \text{LEFT,} & \text{if [case 2]} \\ \text{UP,} & \text{if [case 3]} \end{cases}$$

3. Termination. $F(M, N)$ is the optimal score, and from $Ptr(M, N)$ can trace back optimal alignment



Performance

- Time:
 $O(NM)$
- Space:
 $O(NM)$
- Later we will cover more efficient methods



A variant of the basic algorithm:


- Maybe it is OK to have an unlimited # of gaps in the beginning and end:




- Then, we don't want to penalize gaps in the ends



Different types of overlaps




Example:
2 overlapping “reads” from a sequencing project



Example:
Search for a mouse gene within a human chromosome



The Overlap Detection variant



Changes:

1. Initialization

For all i, j ,
 $F(i, 0) = 0$
 $F(0, j) = 0$

2. Termination

$$F_{OPT} = \max \begin{cases} \max_i F(i, N) \\ \max_j F(M, j) \end{cases}$$



The local alignment problem

Given two strings


$$x = x_1 \dots x_M,$$

$$y = y_1 \dots y_N$$

Find substrings x' , y' whose similarity
(optimal global alignment value)
is maximum

$x = \text{aaaacc} \boxed{\text{cccgggg}} \text{gtta}$


$y = \text{tt} \boxed{\text{cccggg}} \text{aaccaacc}$





Why local alignment

- Genes are shuffled between genomes





Cross-species genome similarity

- 98% of genes are conserved between any two mammals
- >70% average similarity in protein sequence

```
hum_a : GTTGACAATAGAGGGTCTGGCAGAGGCTC----- @ 57331/400001
mus_a : GCTGACAATAGAGGGGCTGGCAGAGGCTC----- @ 78560/400001
rat_a : GCTGACAATAGAGGGGCTGGCAGAGACTC----- @ 112658/369938
fug_a : TTTGTTGATGGGGAGCGTGCATTAATTTTCAGGCTATTGTTAACAGGCTCG @ 36008/68174

hum_a : CTGGCCCGGGTGC GGAGCGTCTGGAGCGGAGCACGCGCTGTCAGCTGGTG @ 57381/400001
mus_a : CTGGCCCCGGTGC GGAGCGTCTGGAGCGGAGCACGCGCTGTCAGCTGGTG @ 78610/400001
rat_a : CTGGCCCCGGTGC GGAGCGTCTGGAGCGGAGCACGCGCTGTCAGCTGGTG @ 112708/369938
fug_a : TGGGCCGAGGTGTTGGATGGCCTGAGTGAAGCACGCGCTGTCAGCTGGCG @ 36058/68174

hum_a : AGCGCACTCTCCTTTTCAGGCAGCTCCCCGGGGAGCTGTGCGGCCACATTT @ 57431/400001
mus_a : AGCGCACTCG-CTTTTCAGGCCGCTCCCCGGGGAGCTGAGCGGCCACATTT @ 78659/400001
rat_a : AGCGCACTCG-CTTTTCAGGCCGCTCCCCGGGGAGCTGCGCGGCCACATTT @ 112757/369938
fug_a : AGCGCTCGCG-----AGTCCCTGCCGTGTCC @ 36084/68174

hum_a : AACACCATCATCACCCCTCCCCGGCCTCCTCAACCTCGGCCTCCTCCTCG @ 57481/400001
mus_a : AACACCGTCGTCA-CCCTCCCCGGCCTCCTCAACCTCGGCCTCCTCCTCG @ 78708/400001
rat_a : AACACCGTCGTCA-CCCTCCCCGGCCTCCTCAACCTCGGCCTCCTCCTCG @ 112806/369938
fug_a : CCGAGGACCCTGA----- @ 36097/68174
```

“atoh” enhancer in human, mouse, rat, fugu fish




The Smith-Waterman algorithm


Idea: Ignore badly aligning regions

Modifications to Needleman-Wunsch:

Initialization: $F(0, j) = F(i, 0) = 0$

Iteration:
$$F(i, j) = \max \begin{cases} 0 \\ F(i - 1, j) - d \\ F(i, j - 1) - d \\ F(i - 1, j - 1) + s(x_i, y_j) \end{cases}$$





The Smith-Waterman algorithm

Termination:

1. If we want the **best** local alignment...


$$F_{\text{OPT}} = \max_{i,j} F(i, j)$$

Find F_{OPT} and trace back

2. If we want **all** local alignments **scoring** $> t$

?? For all i, j find $F(i, j) > t$, and trace back?

Complicated by overlapping local alignments




Waterman–Eggert '87: *find all non-overlapping local alignments with minimal recalculation of the DP matrix*



Scoring the gaps more accurately


Current model:

Gap of length n
incurs penalty $n \times d$




However, gaps usually occur in bunches

Concave gap penalty function $\gamma(n)$
(aka Convex $-\gamma(n)$):



$\gamma(n)$:

$$\text{for all } n, \gamma(n + 1) - \gamma(n) \leq \gamma(n) - \gamma(n - 1)$$



Convex gap dynamic programming

Initialization: same

Iteration:

$$F(i, j) = \max \begin{cases} F(i-1, j-1) + s(x_i, y_j) \\ \max_{k=0 \dots i-1} F(k, j) - \gamma(i-k) \\ \max_{k=0 \dots j-1} F(i, k) - \gamma(j-k) \end{cases}$$

Termination: same

Running Time: $O(N^2M)$ (assume $N > M$)

Space: $O(NM)$



Needleman-Wunsch with affine gaps

Why do we need matrices F, G, H?

Because, perhaps

$$\mathbf{G}(i, j) < \mathbf{V}(i, j)$$

(it is best to align x_i to y_j if we were aligning only $x_1 \dots x_i$ to $y_1 \dots y_j$ and not the rest of x, y),

but on the contrary

$$\mathbf{G}(i, j) - e > \mathbf{V}(i, j) - d$$

(i.e., had we “fixed” our decision that x_i aligns to y_j , we could regret it at the next step when aligning $x_1 \dots x_{i+1}$ to $y_1 \dots y_j$)

Add -d

$$\mathbf{G}(i+1, j) = \mathbf{F}(i, j) - d$$

Add -e

$$\mathbf{G}(i+1, j) = \mathbf{G}(i, j) - e$$



Needleman-Wunsch with affine gaps

Initialization: $V(i, 0) = d + (i - 1) \times e$
 $V(0, j) = d + (j - 1) \times e$

Iteration: $V(i, j) = \max\{ F(i, j), G(i, j), H(i, j) \}$

$$F(i, j) = V(i - 1, j - 1) + s(x_i, y_j)$$

$$G(i, j) = \max \begin{cases} V(i - 1, j) - d \\ G(i - 1, j) - e \end{cases}$$

$$H(i, j) = \max \begin{cases} V(i, j - 1) - d \\ H(i, j - 1) - e \end{cases}$$


Termination: $V(i, j)$ has the best alignment

Time?
Space?



To generalize a bit...

... think of how you would compute optimal alignment with this gap function





Bounded Dynamic Programming

Assume we know that x and y are very similar

Assumption: $\# \text{ gaps}(x, y) < k(N)$


Then, $\begin{matrix} x_i \\ | \\ y_j \end{matrix}$ implies $|i - j| < k(N)$

We can align x and y more efficiently:

Time, Space: $O(N \times k(N)) \ll O(N^2)$



Bounded Dynamic Programming



Initialization:

$F(i,0), F(0,j)$ undefined for $i, j > k$

Iteration:

For $i = 1 \dots M$

For $j = \max(1, i - k) \dots \min(N, i + k)$

$$F(i, j) = \max \begin{cases} F(i - 1, j - 1) + s(x_i, y_j) \\ F(i, j - 1) - d, \text{ if } j > i - k(N) \\ F(i - 1, j) - d, \text{ if } j < i + k(N) \end{cases}$$

Termination: same

Easy to extend to the affine gap case




Outline

- Linear-Space Alignment
- BLAST – local alignment search
- Ultra-fast alignment for (human) genome resequencing



Linear-Space Alignment





Subsequences and Substrings

Definition A string x' is a **substring** of a string x ,
if $x = ux'v$ for some prefix string u and suffix string v

(similarly, $x' = x_i \dots x_j$, for some $1 \leq i \leq j \leq |x|$)

A string x' is a **subsequence** of a string x
if x' can be obtained from x by deleting 0 or more letters

($x' = x_{i_1} \dots x_{i_k}$, for some $1 \leq i_1 \leq \dots \leq i_k \leq |x|$)

Note: a substring is always a subsequence

<u>Example:</u>	$x = \text{abracadabra}$	
	$y = \text{cadabr};$	<i>substring</i>
	$z = \text{brcdbr};$	<i>subsequence, not substring</i>



Hirschberg's algorithm

Given a set of strings x, y, \dots , a **common subsequence** is a string u that is a subsequence of all strings x, y, \dots

- Longest common subsequence

- Given strings $x = x_1 x_2 \dots x_M, y = y_1 y_2 \dots y_N$,
- Find longest common subsequence $u = u_1 \dots u_k$

- Algorithm:

- $$F(i, j) = \max \begin{cases} F(i-1, j) \\ F(i, j-1) \\ F(i-1, j-1) + [1, \text{ if } x_i = y_j; 0 \text{ otherwise}] \end{cases}$$

- $\text{Ptr}(i, j) = (\text{same as in N-W})$


- Termination: trace back from $\text{Ptr}(M, N)$, and prepend a letter to u whenever
 - $\text{Ptr}(i, j) = \text{DIAG}$ **and** $F(i-1, j-1) < F(i, j)$

- Hirschberg's original algorithm solves this in linear space




Introduction: Compute optimal score

It is easy to compute $F(M, N)$ in linear space



```
Allocate ( column[1] )  
Allocate ( column[2] )
```

```
For i = 1....M  
  If i > 1, then:  
    Free( column[ i - 2 ] )  
    Allocate( column[ i ] )  
  For j = 1...N  
    F(i, j) = ...
```



Linear-space alignment

To compute both the optimal score and the optimal alignment:

Divide & Conquer approach:

Notation:

x^r, y^r : reverse of x, y

E.g. $x = \text{accgg}$;

$x^r = \text{ggcca}$


$F^r(i, j)$: optimal score of aligning $x_1^r \dots x_i^r$ & $y_1^r \dots y_j^r$
same as aligning $x_{M-i+1} \dots x_M$ & $y_{N-j+1} \dots y_N$




Linear-space alignment

Lemma: (assume M is even)

$$F(M, N) = \max_{k=0 \dots N} (F(M/2, k) + F^r(M/2, N - k))$$



Example:


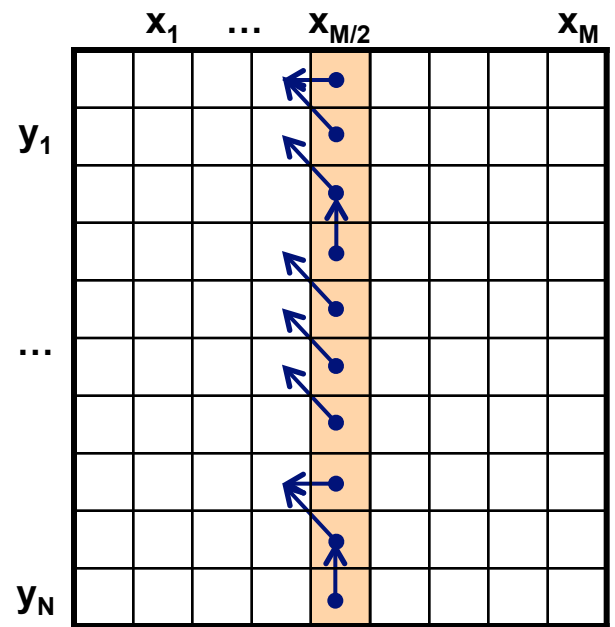




Linear-space alignment

- Now, using 2 columns of space, we can compute for $k = 1 \dots M$, $F(M/2, k)$, $F^r(M/2, N - k)$

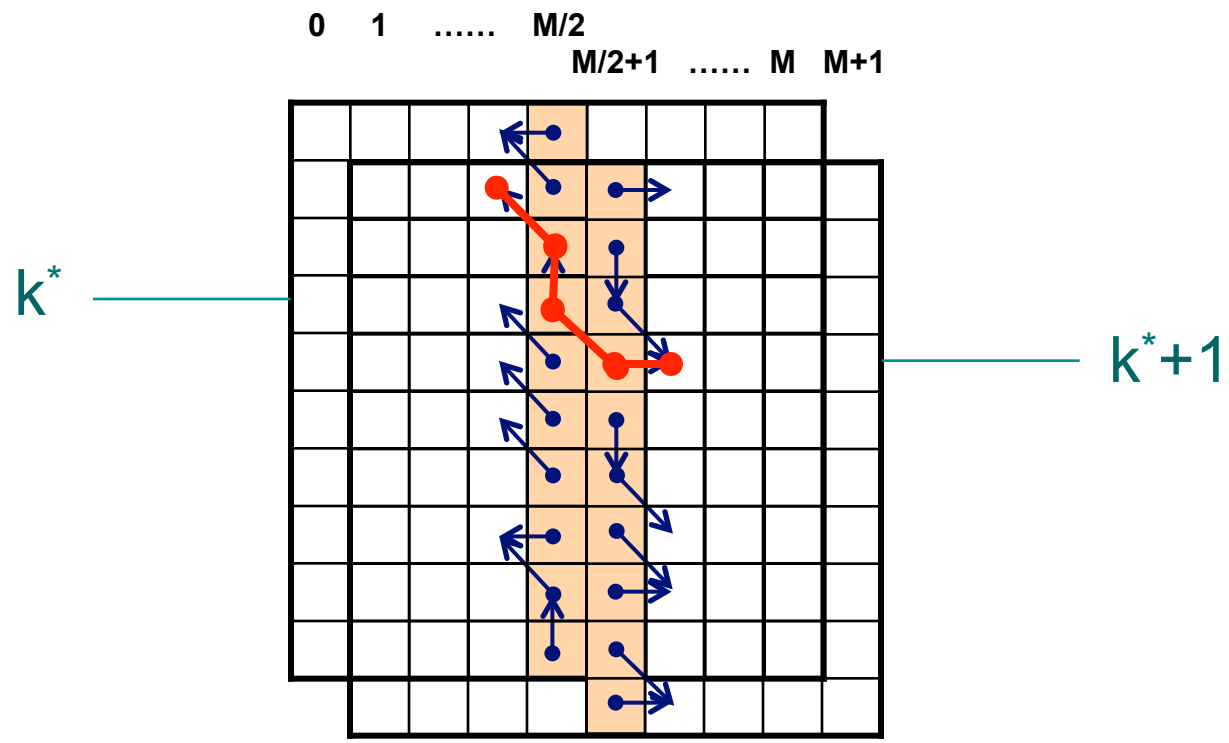
PLUS the backpointers





Linear-space alignment


- Now, we can find k^* maximizing $F(M/2, k) + F^r(M/2, N-k)$
- Also, we can trace the path exiting column $M/2$ from k^*





Linear-space alignment

- Iterate this procedure to the left and right!

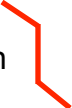




Linear-space alignment

Hirschberg's Linear-space algorithm:

MEMALIGN(l, l', r, r'): (aligns $x_1 \dots x_{l'}$ with $y_r \dots y_{r'}$)

1. Let $h = \lceil (l' - l) / 2 \rceil$
2. Find (in Time $O((l' - l) \times (r' - r))$, Space $O(r' - r)$)
the optimal path, L_h , entering column $h - 1$, exiting column h
Let $k_1 = \text{pos'n at column } h - 2 \text{ where } L_h \text{ enters}$
 $k_2 = \text{pos'n at column } h + 1 \text{ where } L_h \text{ exits}$
3. MEMALIGN($l, h - 2, r, k_1$)
4. Output L_h 
5. MEMALIGN($h + 1, l', k_2, r'$)

Top level call: MEMALIGN(1, M, 1, N)



Linear-space alignment

Time, Space analysis of Hirschberg's algorithm:

To compute optimal path at middle column,

For box of size $M \times N$,

Space: $2N$

Time: cMN , for some constant c

Then, left, right calls cost $c(M/2 \times k^* + M/2 \times (N - k^*)) = cMN/2$

All recursive calls cost

Total Time: $cMN + cMN/2 + cMN/4 + \dots = 2cMN = O(MN)$

Total Space: $O(N)$ for computation,
 $O(N + M)$ to store the optimal alignment