

Pairwise Sequence Alignment

BMI/CS 576
www.biostat.wisc.edu/bmi576/
Mark Craven
craven@biostat.wisc.edu

Pairwise alignment: task definition

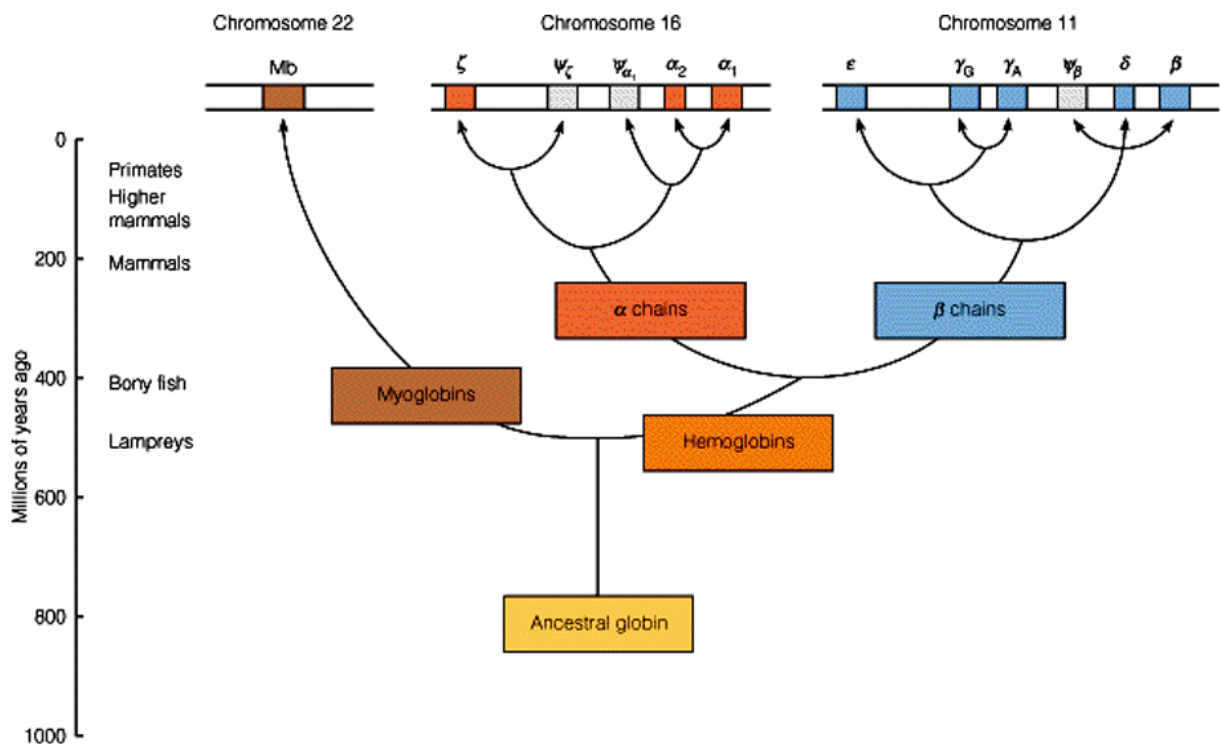
Given

- a pair of sequences (DNA or protein)
- a method for scoring a candidate alignment

Do

- determine the correspondences between substrings in the sequences such that the similarity score is maximized

Homology example: evolution of the globins



Homology

- homologous sequences can be divided into two groups
 - *orthologous sequences*: sequences that differ because they are found in different species (e.g. human α -globin and mouse α -globin)
 - *paralogous sequences*: sequences that differ because of a gene duplication event (e.g. human α -globin and human β -globin, various versions of both)

DNA sequence edits

- substitutions: **ACGA** → **AGGA**
- insertions: **ACGA** → **ACCGGAGA**
- deletions: **ACCGGAGA** → **AGA**
- transpositions: **ACCGGAGA** → **AAGCGGA**
- inversions: **ACCGGAGA** → **ACTCCGA**

Mismatches and gaps

- substitutions in *homologous* sequences result in mismatches in an alignment
- insertions/deletions in *homologous* sequences result in mismatches in an alignment

CA--GATTCGAAT
CGCCGATT---AT


mismatch *gap*

Alignment scales

- for short DNA sequences (gene scale) we will generally only consider
 - substitutions
 - insertions/deletions
- for longer DNA sequences (genome scale) we will consider additional events
 - transpositions
 - inversions
- in this course we will focus on the case of short sequences

Insertions/deletions and protein structure

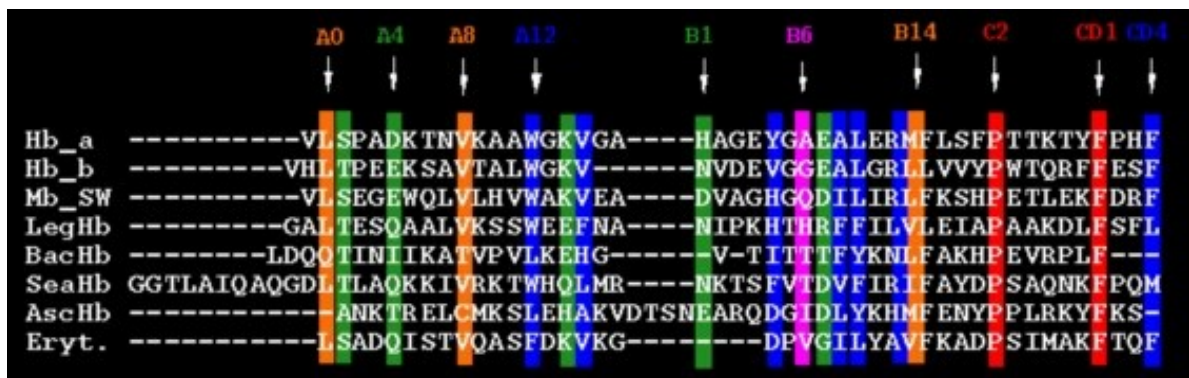

- Why is it that two “similar” sequences may have large insertions/deletions?
 - some insertions and deletions may not significantly affect the structure of a protein



loop structures:
insertions/deletions
here not so significant

Example alignment: globins

- figure at right shows prototypical structure of globins
- figure below shows part of alignment for 8 globins



Issues in sequence alignment

- the sequences we're comparing typically differ in length
- there may be only a relatively small region in the sequences that matches
- we want to allow partial matches (i.e. some amino acid pairs are more substitutable than others)
- variable length regions may have been inserted/deleted from the common ancestral sequence

Types of alignment

- *global*: find best match of both sequences in their entirety
- *local*: find best subsequence match
- *semi-global*: find best match without penalizing gaps on the ends of the alignment

Scoring an alignment: what is needed?

- substitution matrix
 - $s(a,b)$ indicates score of aligning character a with character b
- gap penalty function
 - $w(g)$ indicates cost of a gap of length g

Scoring an alignment

- the score of an alignment is the sum of the scores for pairs of aligned characters plus the scores for gaps
- example: given the following alignment

VAHV---D--DMPNALSALSDLHAHKL

AIQLQVTGVVVTDATLKNLGSVHVSKG

- we would score it by

$$s(\mathbf{V}, \mathbf{A}) + s(\mathbf{A}, \mathbf{I}) + s(\mathbf{H}, \mathbf{Q}) + s(\mathbf{V}, \mathbf{L}) - 3d + s(\mathbf{D}, \mathbf{G}) - 2d$$

...

The space of global alignments

- some possible global alignments for **ELV** and **VIS**

ELV
VIS

-ELV
VIS-

--ELV
VIS--

ELV-
-VIS

E-LV
VIS-

ELV--
--VIS

EL-V
-VIS

- Can we find the highest scoring alignment by enumerating all possible alignments and picking the best?

Number of possible alignments

- given sequences of length m and n
- assume we don't count as distinct $\begin{matrix} C- \\ -G \end{matrix}$ and $\begin{matrix} -C \\ G- \end{matrix}$
- we can have as few as 0 and as many as $\min\{m, n\}$ aligned pairs
- therefore the number of possible alignments is given by

$$\sum_{k=0}^{\min\{m, n\}} \binom{n}{k} \binom{m}{k} = \binom{n+m}{n}$$

Number of possible alignments

- there are

$$\binom{2n}{n} = \frac{(2n)!}{(n!)^2} \approx \frac{2^{2n}}{\sqrt{\pi n}}$$

possible global alignments for 2 sequences of length n


- e.g. two sequences of length 100 have $\approx 10^{77}$ possible alignments
- but we can use *dynamic programming* to find an optimal alignment efficiently

Pairwise alignment via dynamic programming

- first algorithm by Needleman & Wunsch, *Journal of Molecular Biology*, 1970
- *dynamic programming*: solve an instance of a problem by taking advantage of computed solutions for smaller subparts of the problem
- determine best alignment of two sequences by determining best alignment of all prefixes of the sequences

Dynamic programming idea

- consider last step in computing alignment of **AAAC** with **AGC**
- three possible options; in each we'll choose a different pairing for end of alignment, and add this to best alignment of previous characters



Dynamic programming idea

- given an n -character sequence x , and an m -character sequence y
- construct an $(n+1) \times (m+1)$ matrix F
- $F(i, j) =$ score of the best alignment of $x[1\dots i]$ with $y[1\dots j]$

	A	G	C
A			
A			
A			
C			

score of best alignment of AAA to AG

DP algorithm for global alignment with linear gap penalty

- one way to specify the DP is in terms of its recurrence relation:

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

Initializing matrix: global alignment with linear gap penalty

	A	G	C
	0 ← $-d$ ← $-2d$ ← $-3d$		
A	↑ $-d$		
A	↑ $-2d$		
A	↑ $-3d$		
C	↑ $-4d$		

DP algorithm sketch: global alignment

- initialize first row and column of matrix
- fill in rest of matrix from top to bottom, left to right
- for each $F(i, j)$, save pointer(s) to cell(s) that resulted in best score
- $F(m, n)$ holds the optimal alignment score; trace pointers back from $F(m, n)$ to $F(0, 0)$ to recover alignment

Global alignment example

- suppose we choose the following scoring scheme:


$$s(x_i, y_i) =$$

+1 when $x_i = y_i$

-1 when $x_i \neq y_i$

d (penalty for aligning with a gap) = 2

Global alignment example



one optimal alignment



x: A A A C
y: A G - C

DP comments

- works for either DNA or protein sequences, although the substitution matrices used differ
- finds an optimal alignment
- the exact algorithm (and computational complexity) depends on gap penalty function (we'll come back to this issue)

Equally optimal alignments

- many optimal alignments may exist for a given pair of sequences
- can use preference ordering over paths when doing traceback



- *highroad* and *lowroad* alignments show the two most different optimal alignments

Highroad & lowroad alignments

		A	G	C	
		0	-2	-4	-6
A	-2	1	-1	-3	
A	-4	-1	0	-2	
A	-6	-3	-2	-1	
C	-8	-5	-4	-1	

highroad alignment

x: A A A C
y: A G - C

lowroad alignment

x: A A A C
y: - A G C

Computational complexity

- initialization: $O(m)$, $O(n)$ where sequence lengths are m , n
- filling in rest of matrix: $O(mn)$
- traceback: $O(m + n)$
- hence, if sequences have nearly same length, the computational complexity is

$$O(n^2)$$

Local alignment

- so far we have discussed *global alignment*, where we are looking for best match between sequences from one end to the other
- often we want a *local alignment*, the best match between subsequences of x and y

Example local alignment

- aligning my name against the sequence for dTDP-4-dehydrorhamnose reductase from the bacterium *opitutus terrae*

...LSGAYHLAASGHTSWHGFASAIIDLMPLDARKCRAVEAIT...
MARKCRAVEN

Local alignment motivation

- useful for comparing protein sequences that share a common *motif* (conserved pattern) or *domain* (independently folded unit) but differ elsewhere
- useful for comparing DNA sequences that share a similar *motif* but differ elsewhere
- useful for comparing protein sequences against *genomic DNA sequences* (long stretches of uncharacterized sequence)
- more sensitive when comparing highly diverged sequences

Local alignment DP algorithm

- original formulation: Smith & Waterman, *Journal of Molecular Biology*, 1981
- interpretation of array values is somewhat different:
 $F(i, j)$ = score of the best alignment of a suffix of $x[1..i]$ and a suffix of $y[1..j]$

Local alignment DP algorithm

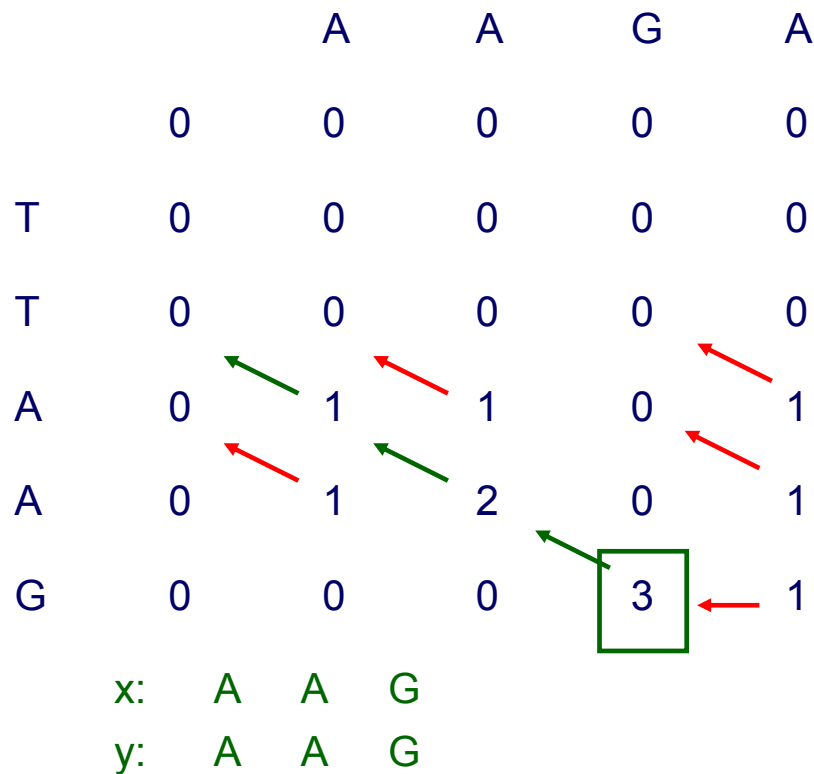
- the recurrence relation is slightly different than for global algorithm

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

Local alignment DP algorithm

- initialization: first row and first column initialized with 0's
- traceback:
 - find maximum value of $F(i, j)$; can be anywhere in matrix
 - stop when we get to a cell with value 0

Local alignment example



More on gap penalty functions

- a gap of length k is more probable than k gaps of length 1
 - a gap may be due to a single mutational event that inserted/deleted a stretch of characters
 - separated gaps are probably due to distinct mutational events
- a linear gap penalty function treats these cases the same
- it is more common to use gap penalty functions involving two terms
 - a penalty d associated with opening a gap
 - a smaller penalty e for extending the gap

Gap penalty functions

linear

$$w(g) = -g \times d$$

affine

$$w(g) = \begin{cases} -d - (g - 1)e, & g \geq 1 \\ 0, & g = 0 \end{cases}$$

Dynamic programming for the affine gap penalty case

- to do in $O(n^2)$ time, need 3 matrices instead of 1

$M(i, j)$ best score given that $x[i]$ is aligned to $y[j]$

$I_x(i, j)$ best score given that $x[i]$ is aligned to a gap

$I_y(i, j)$ best score given that $y[j]$ is aligned to a gap

Global alignment DP for the affine gap penalty case

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

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

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

Global alignment DP for the affine gap penalty case

- initialization

$$M(0, 0) = 0$$

$$I_x(i, 0) = -d - (i-1)e \quad \text{for } i > 0$$

$$I_y(0, j) = -d - (j-1)e \quad \text{for } j > 0$$

other cells in top row and leftmost column = $-\infty$

- traceback

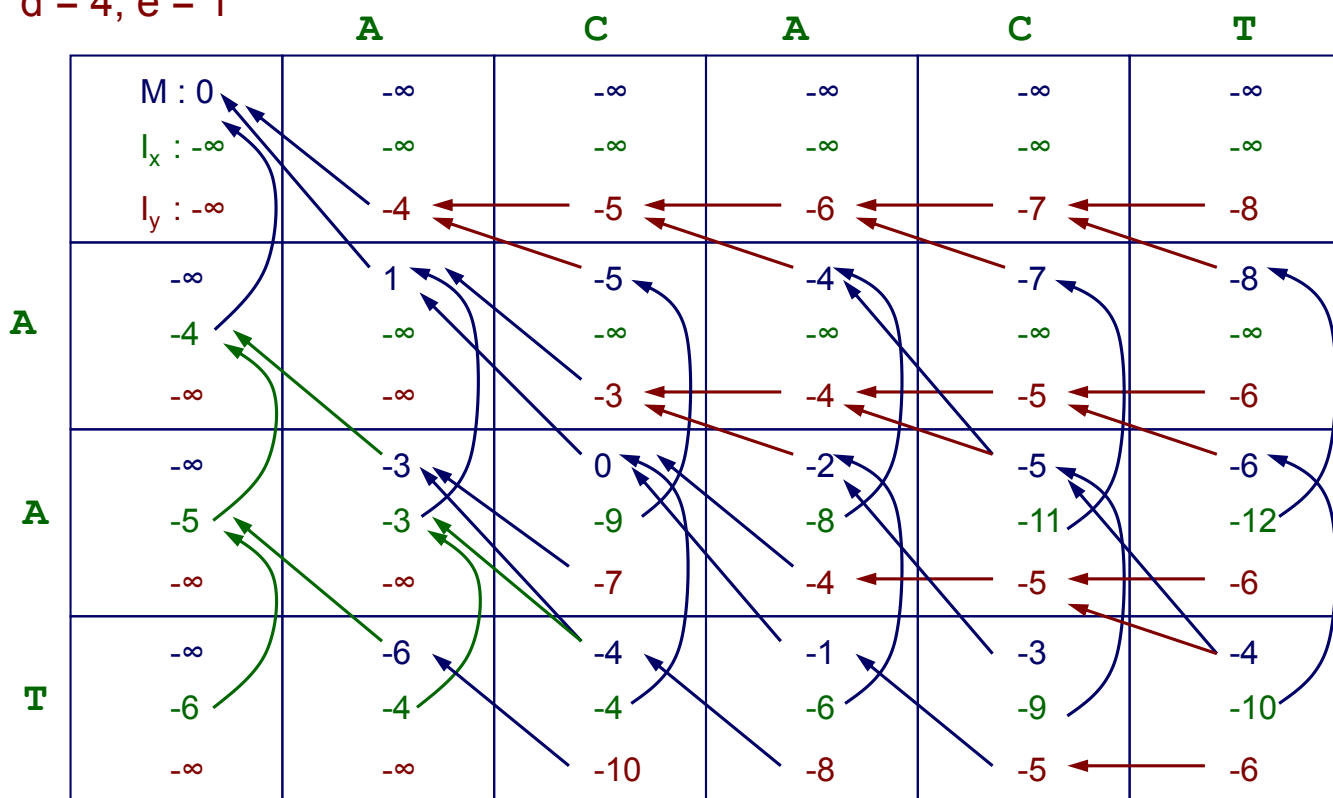
- start at largest of $M(m, n)$, $I_x(m, n)$, $I_y(m, n)$

- stop at $M(0, 0)$

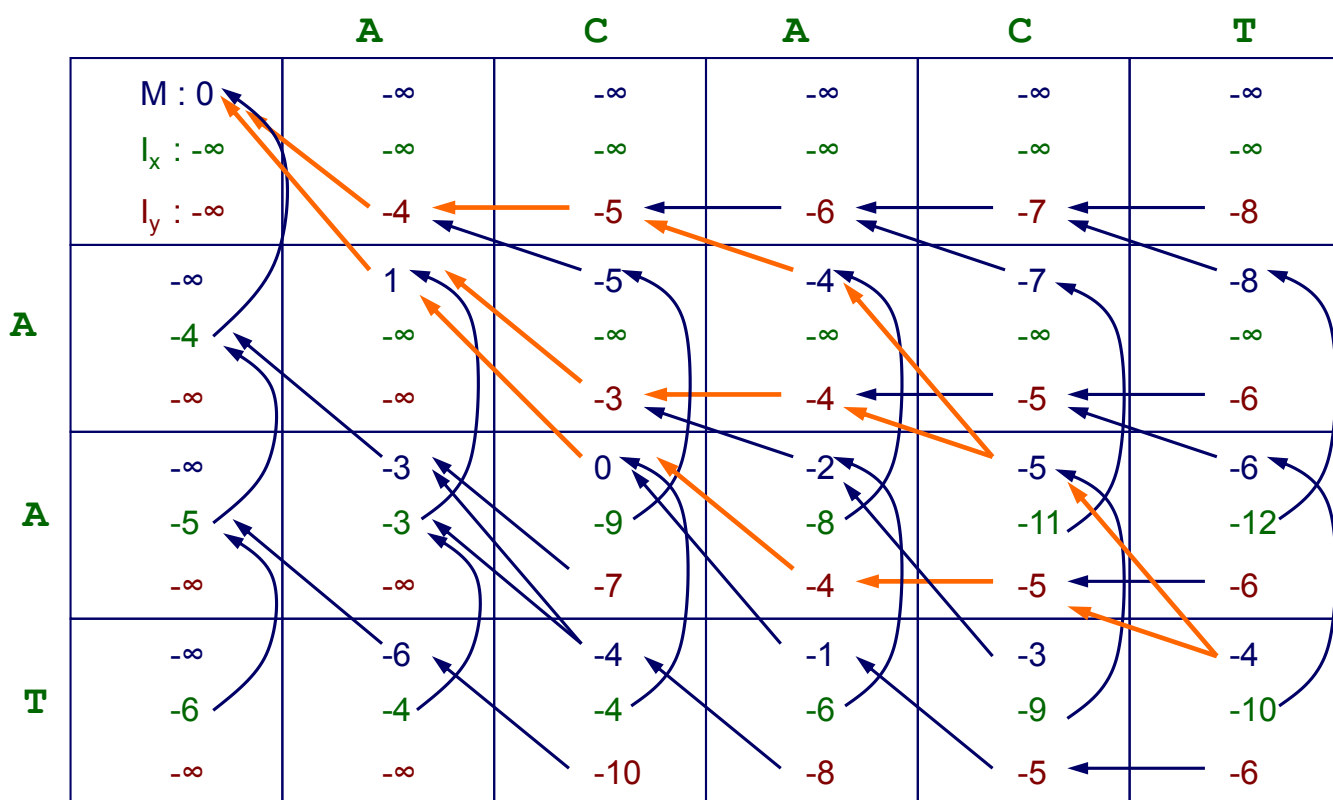
- note that pointers may traverse all three matrices

Global alignment example (affine gap penalty)

$d = 4, e = 1$



Global alignment example (continued)



three optimal alignments: **ACACT** **ACACT** **ACACT**
AA--T **A--AT** **--AAT**

Why three matrices are needed

- consider aligning the sequences **WFP** and **FW** using $d = 5$, $e = 1$ and the following values from the BLOSUM-62 substitution matrix:

$$\begin{aligned} s(\mathbf{F}, \mathbf{W}) &= 1 & s(\mathbf{W}, \mathbf{W}) &= 11 \\ s(\mathbf{F}, \mathbf{F}) &= 6 & s(\mathbf{W}, \mathbf{P}) &= -4 \\ s(\mathbf{F}, \mathbf{P}) &= -4 \end{aligned}$$

- the matrix shows the highest-scoring partial alignment for each pair of prefixes

	W	F	P
	0	-5	-6
F	-5	1	1
W	-6	6	2

Annotations:

- A bracket on the right side of the matrix points to the cell (F, F) and (W, F) with the text: **-WFP** and **FW--** optimal alignment
- A bracket on the bottom side of the matrix points to the cell (F, F) and (W, F) with the text: **WF** and **FW** best alignment of these prefixes; to get optimal alignment, need to also remember
- An arrow points from the text "need to also remember" to a bracket on the right side of the matrix pointing to the cell (F, F) and (W, F) with the text: **-WF** and **FW-**

Local alignment DP for the affine gap penalty case

$$M(i, j) = \max \begin{cases} M(i-1, j-1) + s(x_i, y_j) \\ I_x(i-1, j-1) + s(x_i, y_j) \\ I_y(i-1, j-1) + s(x_i, y_j) \\ \mathbf{0} \end{cases}$$

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

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

Local alignment DP for the affine gap penalty case

- initialization

$$M(0,0) = 0$$

$$M(i,0) = 0$$

$$M(0,j) = 0$$

cells in top row and leftmost column of $I_x, I_y = -\infty$

- traceback

- start at largest $M(i,j)$

- stop at $M(i,j) = 0$

Gap penalty functions

- linear: $w(g) = -g \times d$

- affine:
$$w(g) = \begin{cases} -d - (g-1)e, & g \geq 1 \\ 0, & g = 0 \end{cases}$$

- convex: as gap length increases, magnitude of penalty for each additional character decreases

e.g. $w(g) = -d - \log(g) \times e$

Computational complexity and gap penalty functions

linear: $O(n^2)$

affine: $O(n^2)$

convex: $O(n^2 \log n)$

general: $O(n^3)$

* assuming two sequences of length n

Alignment (global) with general gap penalty function

why the general case has time complexity $O(n^3)$

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

k ranges over previous coordinates

consider every previous element in the row

consider every previous element in the column

Pairwise alignment summary

- the number of possible alignments is exponential in the length of sequences being aligned
- dynamic programming can find optimal-scoring alignments in polynomial time
- the specifics of the DP depend on
 - local vs. global alignment
 - gap penalty function
- affine penalty functions are most commonly used