

Practical considerations of working with sequencing data

File Types

- Fastq -> aligner -> reference(genome) coordinates
- Coordinate files
 - SAM/BAM – most complete, contains all of the info in fastq and more!
 - Bedgraph – read density along the genome
 - Bed file – Read density reported in large continuous intervals
 - Genes/transcript and transcript structure
 - Transcription factor binding regions
- If someone does a sequencing experiment usually one of these is available and deposited in a public database

SAM/BAM

(1) The query name of the read is given (M01121...)

(2) The flag value is 163 (this equals 1+2+32+128)

(3) The reference sequence name, chrM, refers to the mitochondrial genome

(4) Position 480 is the left-most coordinate position of this read

(5) The Phred-scaled mapping quality is 60 (an error rate of 1 in 10⁶)

(6) The CIGAR string (148M2S) shows 148 matches and 2 soft-clipped (unaligned) bases

```
home/bioinformatics$ samtools view 030c_S7.bam | less
M01121:5:000000000-A2DTN:1:2111:20172:15571      163      chrM
480      60      148M2S      =      524      195      AATCTCATCAAT
ACAACCCTCGCCCATCCTACCCAGCACACACACACCGCTGCTAACCCCATACCCCGAACC
AACCAAACCCCAAAGACACCCCCACAGTTTATGTAGCTTACCTCCTCAAAGCAATAACC
TGAAAATGTTTAGACGGG      BBBBFB5@FFGGGFEGGGEGAAACGHFHFEGGAGFFH
AEFDGG?E?EGGGFGHFGHF?FFCHF00E@EGFGGEE1FFEEHGBEFFFGGG@</0
1BG212222>F21@F11FGFG1@1?GC<G11?1?FGDGGF=GHFFFHC.-
RG:Z:Sample7      XC:i:148      XT:A:U      NM:i:3      SM:i:37
AM:i:37      X0:i:1      X1:i:0      XM:i:3      XO:i:0      XG:i:0      MD:Z:19C109C0A17
```

(7) An = sign shows that the mate reference matches the reference name

(8) The 1-based left position is 524

(9) The insert size is 195 bases

(10) The sequence begins AATCT and ends ACGGG (its length is 150 bases)

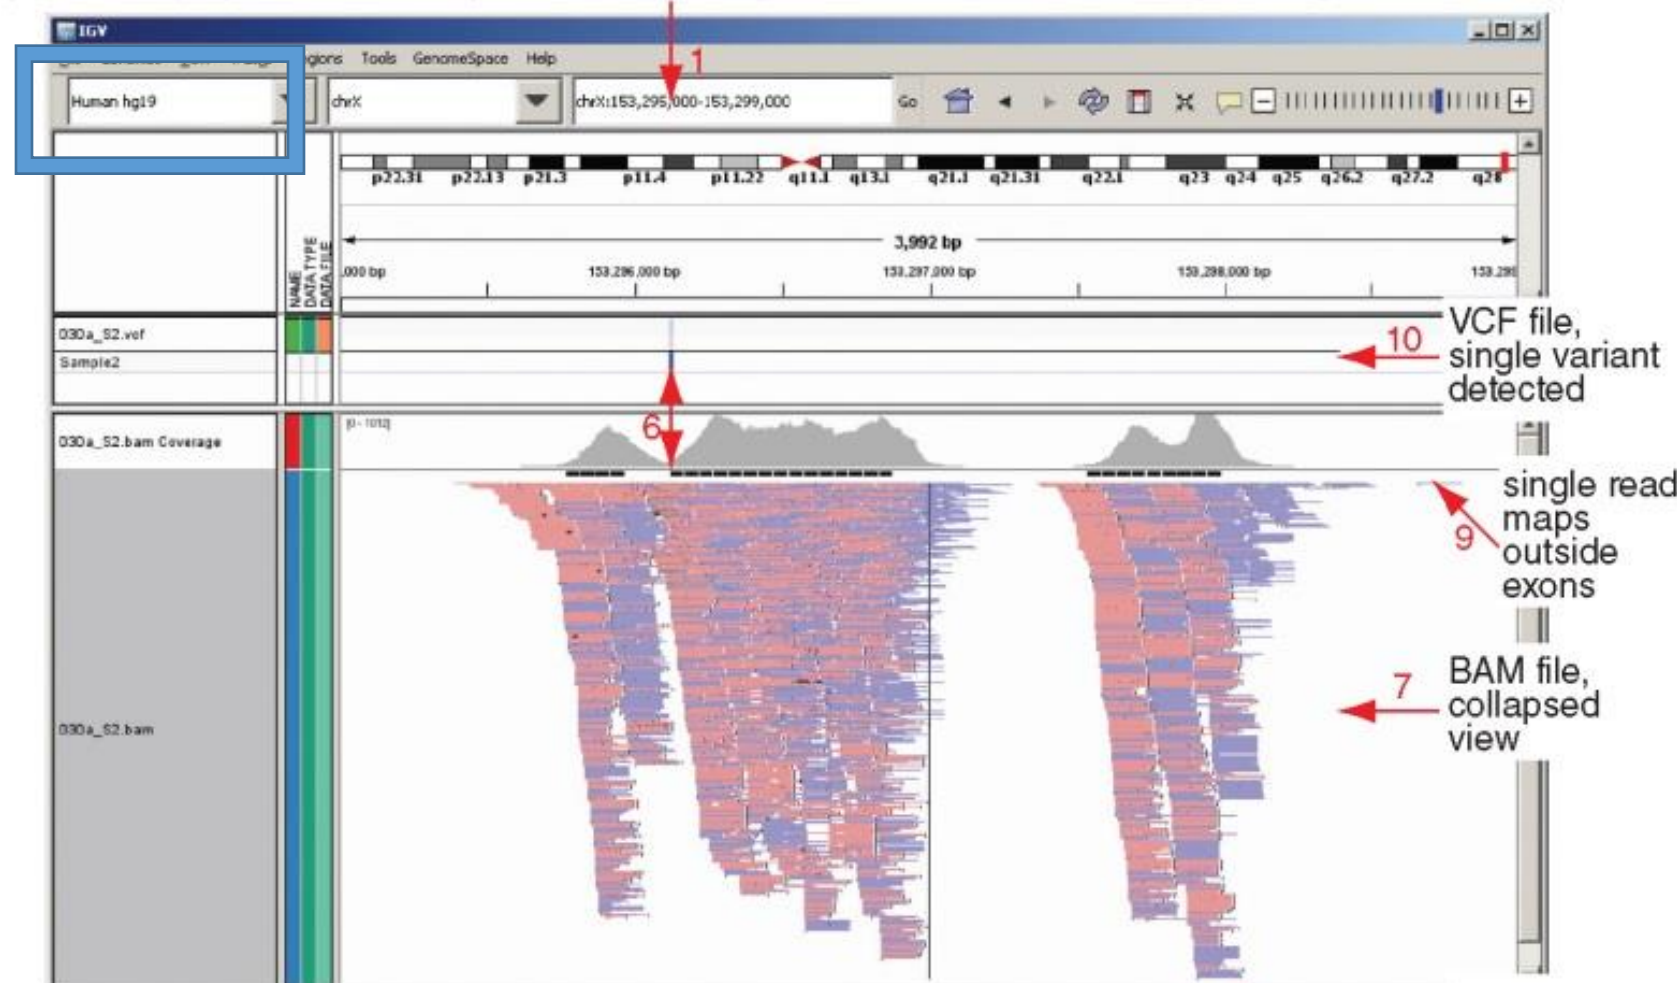
(11) Each base is assigned a quality score (from BBBB ending FHC.-)

(12) This read has additional, optional fields that accompany the MiSeq analysis

Viewing genome coordinate files with IGV

- Integrated Genome Browser
- Cross-platform application
- Knows about common genomes
- Genome version is important!

(a) IGV display of a BAM file (at two resolutions) and a VCF in the *MECP2* gene region



Different assemblies

- Genome coordinates different between genome assemblies
 - Differences accumulate over chromosome length
- You have to know which assembly was used
- Sequencing files are non-randomly distributed relative to genes
 - RNAseq—should align with exons
 - TF binding sites—biased towards promoter regions

Human

- Source: UCSC Genome Bioinformatics, <http://genome.ucsc.edu/>
- Assemblies:
 - UCSC hg19 (GCA_000001405.1), February 2009
 - UCSC hg18 (NCBI build 36.1), March 2006
 - UCSC hg17 (NCBI build 35), May 2004
 - UCSC hg16 (NCBI build 34), July 2003

Human: 1000 Genomes

- Source: 1000 Genomes, <http://www.1000genomes.org/>
- Assembly: b37, October 2009
- Assembly: b36 (1kg ref), December 2008



Mouse (*Mus musculus*)

- Source: UCSC Genome Bioinformatics, <http://genome.ucsc.edu/>
- Assemblies:
 - UCSC mm9 (NCBI build 37), July 2007
 - UCSC mm8 (NCBI build 36), February 2006
 - UCSC mm7 (NCBI build 35), August 2005

Converting coordinates

- UCSC liftOver -- converts genome coordinates
- Convert from one assembly to another
- Cross organism conversion
 - Mammals/vertebrates

mm10ToLoxAfr3.over.chain.gz	20-Mar-2012	15:38	51M
mm10ToMacEug2.over.chain.gz	24-Mar-2012	11:54	12M
mm10ToMelGal1.over.chain.gz	03-Apr-2012	11:53	7.0M
mm10ToMelUnd1.over.chain.gz	30-Mar-2012	04:25	7.1M
mm10ToMicMur1.over.chain.gz	13-Mar-2012	22:10	55M
mm10ToMm9.over.chain.gz	30-Apr-2012	21:52	940K
mm10ToMonDom5.over.chain.gz	30-Mar-2012	19:24	20M
mm10ToMyoLuc2.over.chain.gz	22-Mar-2012	09:03	49M
mm10ToNomLeu1.over.chain.gz	08-Mar-2012	22:47	66M
mm10ToNomLeu2.over.chain.gz	14-Apr-2012	21:16	65M
mm10ToOchPri2.over.chain.gz	24-Mar-2012	06:12	33M



Sequence Alignment

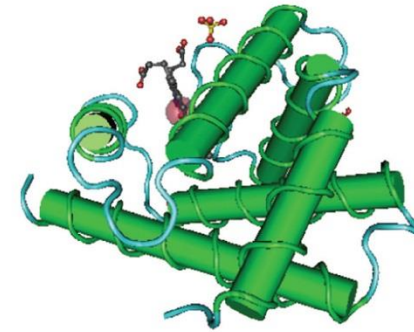
To do:

- Global alignment
- Local alignment
- Scoring
 - Gaps
 - Scoring matrices
- Database Search
 - Statistical Significance
- Multiple Sequence alignment


Why compare sequences

- Given a new sequence, infer its function based on similarity to another sequence
- Find important molecular regions – conserved across species
- Determine 3d structure with homology modeling
- **Homologs**-sequences that descended from a common ancestral sequence
 - **Orthologs**- separated by speciation
 - **Paralogs** separated by duplication in a single genome
- Basic unit of protein homology is a sufficient functional unit—typically much smaller than a whole gene


(a) Human myoglobin (3RGK)






(b) Human hemoglobin tetramer (2H35)



(c) Human beta globin (subunit of 2H35)



(d) Pairwise alignment of beta globin and myoglobin



DNA vs Protein alignments

- Protein coding
 - Typically compared in amino acid space
 - Amino acid change slower than nucleotides
 - Some nucleotides can change without any change to a.a. sequence
 - Different levels of amino acid similarity can be accounted for
 - Not all a.a. changes are equally disruptive
 - Can detect very remote homology
- Non-coding regions
 - Smaller alphabet requires more matches to achieve significance
 - No notion of similarity—match or no match
 - Diverge more rapidly though some are very conserved at short evolutionary distances

What is a good sequence alignment

- Theory: If two sequences are homologous we want to match up the residues such that each residue is descendant from a common ancestral residue
- Practice: approximate string matching
 - introduce gaps and padding to find best matching between two strings



AGGCTATCACCTGACCTCCAGGCCGATGCCC
 TAGCTATCACGACCGCGGTCGATTTGCCCGAC





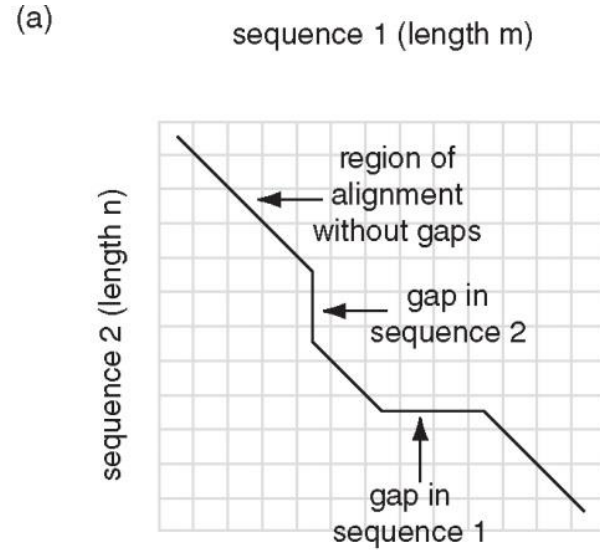


-AGGCTATCACCTGACCTCCAGGCCGA--TGCCC---
 TAG-CTATCAC--GACCGC--GGTCGATTTGCCCGAC

Efficient alignment

- What is the best alignment? – we need a scoring metric
 - Basic scoring metric (1 for matching, 0 for mismatching, 0 for a gap)
- Number of possible alignments is exponential in string length
- Scoring is local
- we apply **dynamic programming**
- **dynamic programming** –solve a large problem in terms of smaller subproblems
- Requirements
 - There is 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

Matrix representation of an alignment



Dynamical programming approach

- Score the optimal alignment up to every (i,j) $F(i,j)$
- Scoring is local so $F(i,j)$ depends only 3 other values

(a)


		Sequence 2								
		F	M	D	T	P	L	N	E	
Sequence 1		0	-2	-4	-6	-8	-10	-12	-14	-16
	F	-2								
	K	-4								
	H	-6								
	M	-8								
	E	-10								
	D	-12								
	P	-14								
	L	-16								
	E	-18								

(b)

$$\text{Score} = \text{Max} \begin{cases} F(i-1, j-1) + s(x_i, y_i) \\ F(i-1, j) - \text{gap penalty} \\ F(i, j-1) - \text{gap penalty} \end{cases}$$

Score (this example) = +1 (match)
 -2 (mismatch)
 -2 (gap penalty)

Global alignment



$$F_{i,j} = \text{MAX}$$

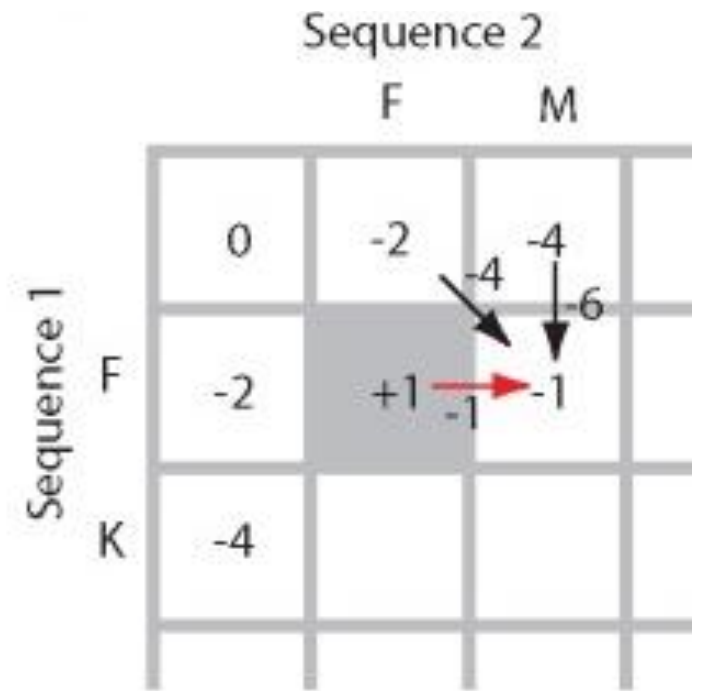
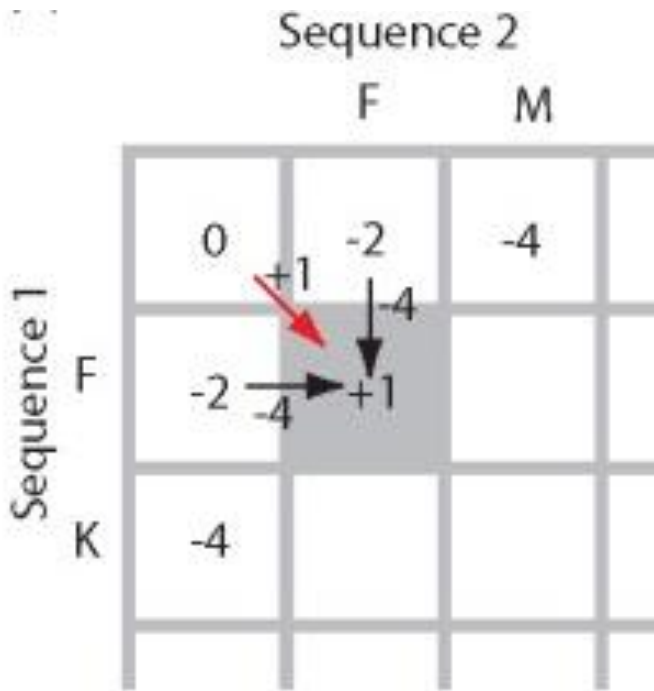
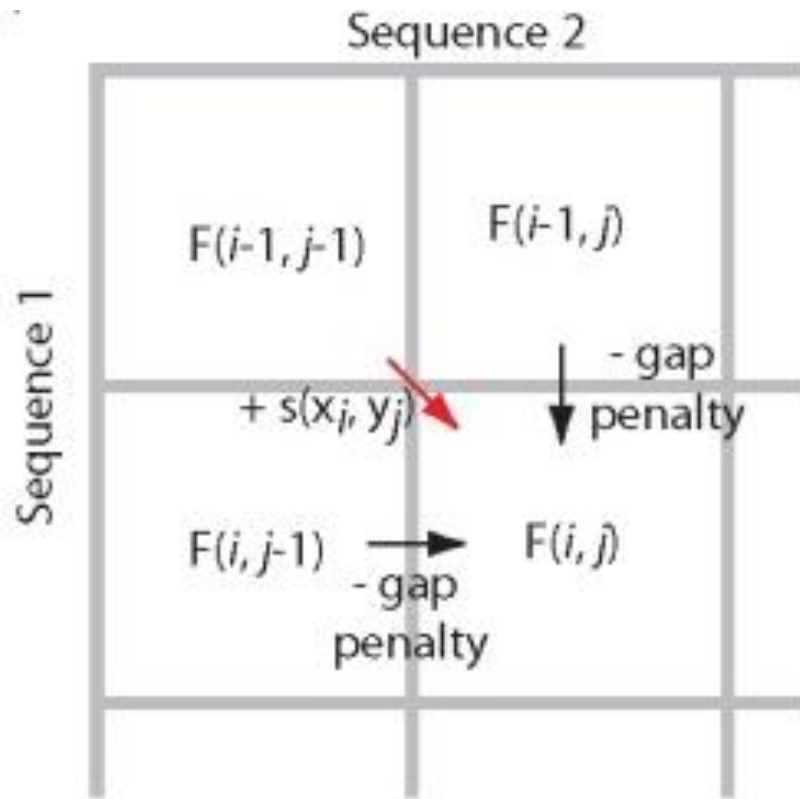
$$\left\{ \begin{array}{l} F_{i-1,j-1} + \text{Score}(S_i, T_j) \\ F_{i,j-1} + \text{gp} \\ F_{i-1,j} + \text{gp} \end{array} \right.$$

Gap penalty

Needleman & Wunsch, 1970

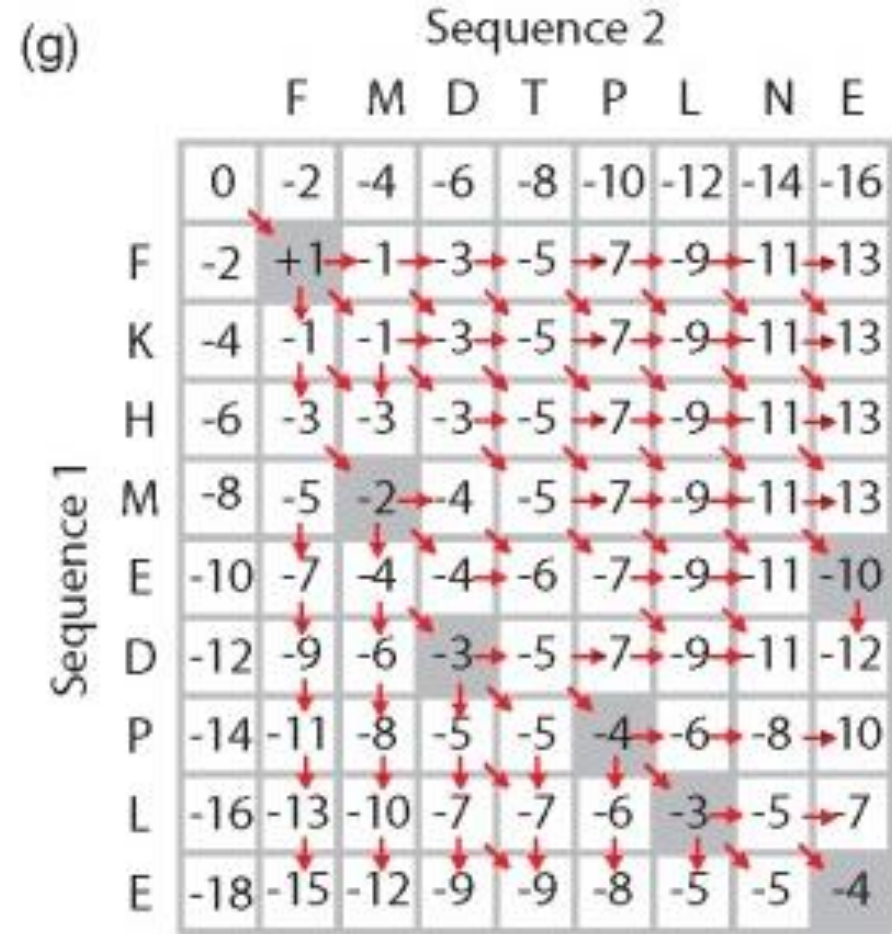

Example

Score (this example) = +1 (match)
-2 (mismatch)
-2 (gap penalty)




Keep track of the argmax!

Matrix filled out

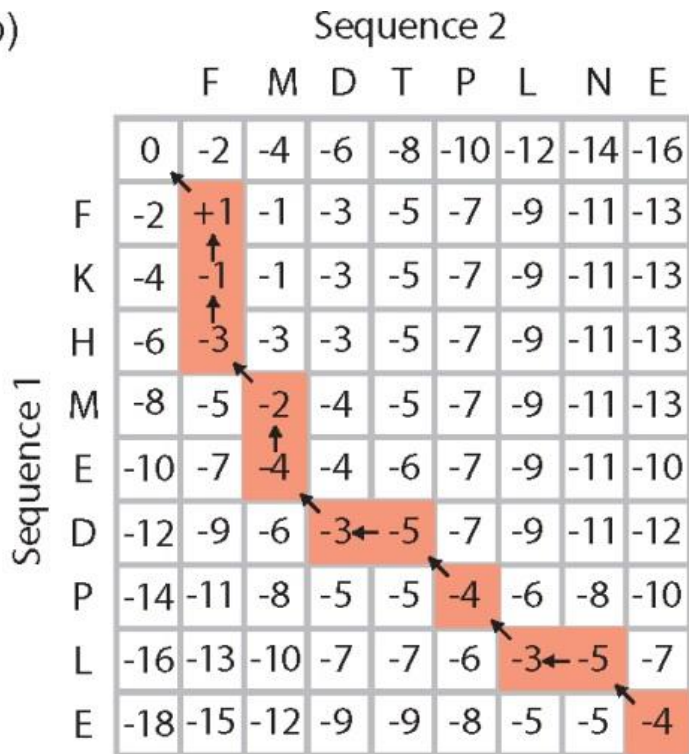


Finding the optimal alignment


(a)



(b)



(c)



Complete Algorithm

- Initialization.

$$F(0,0) = 0$$

$$F(0, j) = -j \times g_o$$

$$F(i, 0) = -i \times g_o$$

- Main Iteration. Filling-in partial alignments

For each $i=1.....M$

For each $j = 1.....N$

$$F(i, j) = \max(F(i-1,j-1)+s(x_i, y_j)...$$

$$F(i-1, j) - g_p,...$$

$$F(i, j-1) - g_p)$$


$$Ptr(i,j) = \text{DIAG}$$

LEFT

UP

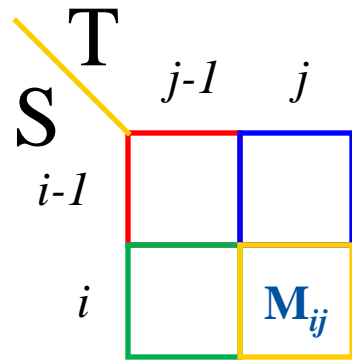
Local alignment

- Given two sequences, S and T, find two subsequences, s and t, whose alignment has the highest “score” amongst all subsequence pairs.
- Two genes in different species may be similar over short conserved regions and dissimilar over remaining regions.
- Example:
 - Homeobox genes have a short region called the *homeodomain* that is highly conserved between species.
 - A global alignment would not find the homeodomain because it would try to align the ENTIRE sequence
- Genes can have local similarity because of variable domain composition



EGR4_HUMAN	KA	[FACPVE	SCVRS	FARSD	ELNRHL	RIRH]	TGHKP	[FQCRIC	LRNFS	RSDDL	TSHV	RTH]	TGEKP	[FACDV	--CG	RRFAR	SDEK	KRH	SKVH]	
EGR4_RAT	KA	[FACPVE	SCVRT	FARSD	ELNRHL	RIRH]	TGHKP	[FQCRIC	LRNFS	RSDDL	TTHV	RTH]	TGEKP	[FACDV	--CG	RRFAR	SDEK	KRH	SKVH]	
EGR3_HUMAN	RP	[HACPAE	GCDRR	FRSRS	DELTRH	LRIH]	TGHKP	[FQCRIC	MRSFS	RSDDL	TTHIR	TTH]	TGEKP	[FACEF	--CG	RKFAR	SDEK	RHAK	IAH]	
EGR3_RAT	RP	[HACPAE	GCDRR	FRSRS	DELTRH	LRIH]	TGHKP	[FQCRIC	MRSFS	RSDDL	TTHIR	TTH]	TGEKP	[FACEF	--CG	RKFAR	SDEK	RHAK	IAH]	
EGR1_HUMAN	RP	[YACPVE	SCDRR	FRSRS	DELTRH	IRIH]	TGQKP	[FQCRIC	MNFS	RSDDL	TTHIR	TTH]	TGEKP	[FACDI	--CG	RKFAR	SDEK	RHTK	IAH]	
EGR1_MOUSE	RP	[YACPVE	SCDRR	FRSRS	DELTRH	IRIH]	TGQKP	[FQCRIC	MNFS	RSDDL	TTHIR	TTH]	TGEKP	[FACDI	--CG	RKFAR	SDEK	RHTK	IAH]	
EGR1_RAT	RP	[YACPVE	SCDRR	FRSRS	DELTRH	IRIH]	TGQKP	[FQCRIC	MNFS	RSDDL	TTHIR	TTH]	TGEKP	[FACDI	--CG	RKFAR	SDEK	RHTK	IAH]	
EGR1_BRARE	RP	[YACPVE	TCDRR	FRSRS	DELTRH	IRIH]	TGQKP	[FQCRIC	MNFS	RSDDL	TTHIR	TTH]	TGEKP	[FACDI	--CG	RKFAR	SDEK	RHTK	IAH]	
EGR2_RAT	RP	[YPCPAE	GCDRR	FRSRS	DELTRH	IRIH]	TGHKP	[FQCRIC	MNFS	RSDDL	TTHIR	TTH]	TGEKP	[FACDY	--CG	RKFAR	SDEK	RHTK	IAH]	
EGR2_XENLA	RP	[YPCPAE	GCDRR	FRSRS	DELTRH	IRIH]	TGHKP	[FQCRIC	MNFS	RSDDL	TTHIR	TTH]	TGEKP	[FACDY	--CG	RKFAR	SDEK	RHTK	IAH]	
EGR2_MOUSE	RP	[YPCPAE	GCDRR	FRSRS	DELTRH	IRIH]	TGHKP	[FQCRIC	MNFS	RSDDL	TTHIR	TTH]	TGEKP	[FACDY	--CG	RKFAR	SDEK	RHTK	IAH]	
EGR2_HUMAN	RP	[YPCPAE	GCDRR	FRSRS	DELTRH	IRIH]	TGHKP	[FQCRIC	MNFS	RSDDL	TTHIR	TTH]	TGEKP	[FACDY	--CG	RKFAR	SDEK	RHTK	IAH]	
EGR2_BRARE	RP	[YPCPAE	GCDRR	FRSRS	DELTRH	IRIH]	TGHKP	[FQCRIC	MNFS	RSDDL	TTHIR	TTH]	TGEKP	[FACDF	--CG	RKFAR	SDEK	RHTK	IAH]	
MIG1_KLULA	--	[-----	-----	-----	-----	-----]	---RP	[YVCPIC	QGRGF	HRL	EHQ	TRH	IRTH]	TGERP	[HACDF	PGC	KRFS	RSDEL	TRH	RRIRH]
MIG1_KLUMA	--	[-----	-----	-----	-----	-----]	---RP	[YMCPI	CHRGF	HRL	EHQ	TRH	IRTH]	TGERP	[HACDF	PGC	AKRFS	RSDEL	TRH	RRIRH]
MIG1_YEAST	--	[-----	-----	-----	-----	-----]	---RP	[HACPI	CHRAF	HRL	EHQ	TRH	MRIRH]	TGEKP	[HACDF	PGC	VKRFS	RSDEL	TRH	RRIRH]
MIG2_YEAST	--	[-----	-----	-----	-----	-----]	---RP	[FRCDT	CHRGF	HRL	EHK	KRH	LRTH]	TGEKP	[HCAF	PGC	GKFS	RSDEL	KR	HMRTH]
		[]	.*	[. * * * * *	.*	. * * *]	***.*	[. * * * * *	. * * *]	.*	.*	

Local alignment



$$F_{i,j} = \text{MAX}$$

0

No pointer, we start over

$$F_{i-1,j-1} + \text{Score}(S_i, T_j)$$

$$F_{i,j-1} + \text{go}$$

$$F_{i-1,j} + \text{go}$$

Gap penalty

Smith & Waterman, 1981

Similarity Scoring Expected value:
negative for random alignments
positive for highly similar sequences

Local alignment

- Initialization

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

- Iteration

for $i=1,\dots,M$

for $j=1,\dots,N$

- calculate optimal $F(i,j)$

- store $\text{Ptr}(i,j)$ if score is positive


- Termination

Find the end of the best alignment with $F_{\text{OPT}} = \max\{i,j\} F(i,j)$ and trace back

OR

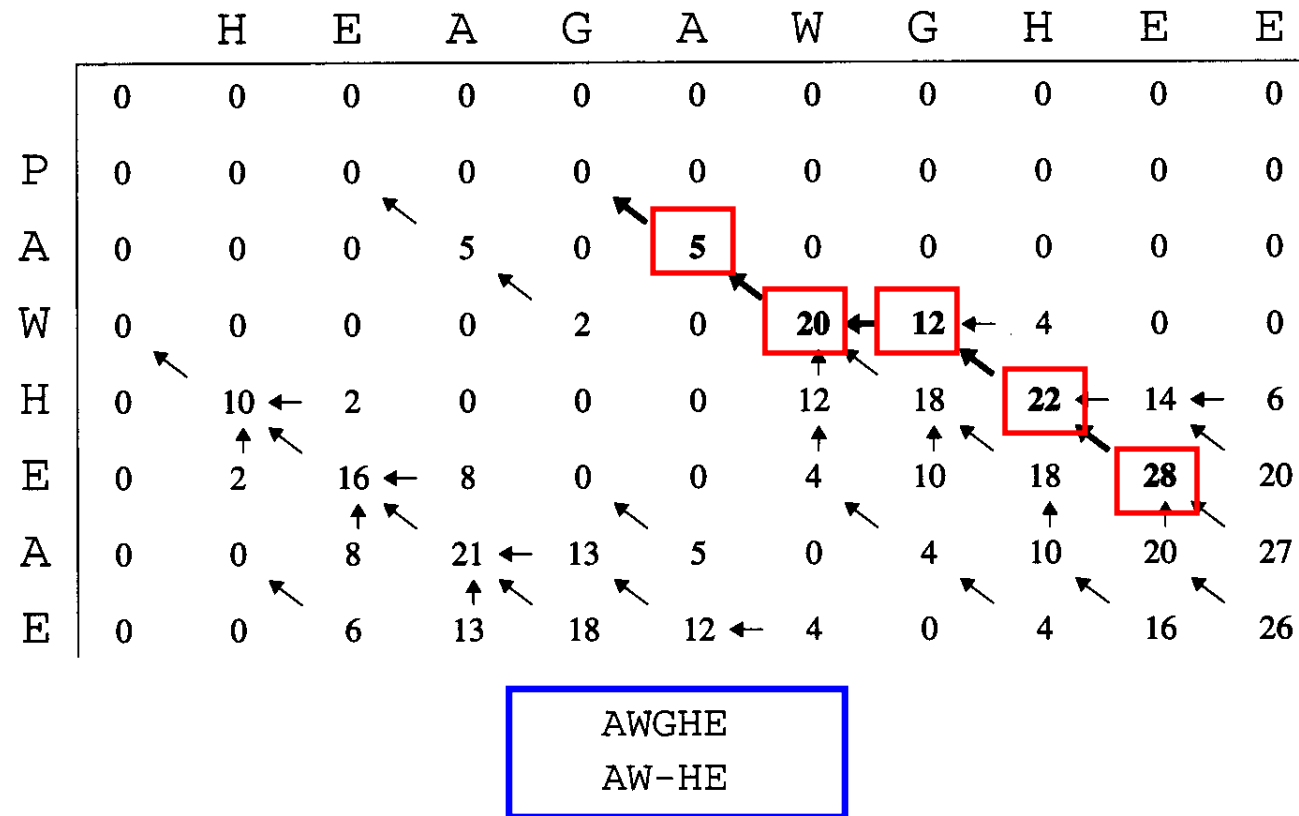
Find all alignments with $F(i,j) > \text{threshold}$ and trace back

Local vs. global alignment






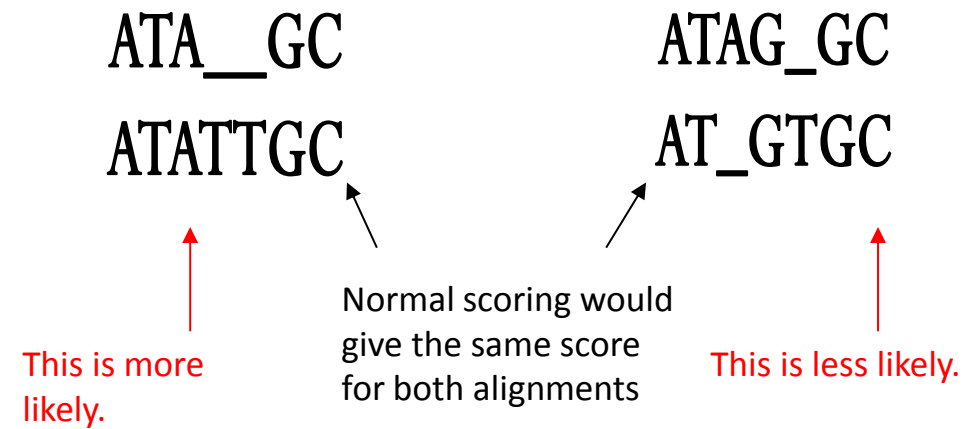
HEAGAWGHEE-E
 --P-AW-HEAE

Local vs. global alignment (cntd)



More accurate gap model


- In nature, gaps often come as a single event rather than a series of single gaps
- Linear gap penalty is too stringent
- Convex gap penalty is expensive
 - Have to keep track of the length of gaps
- Compromise – Affine gap penalty
 - $\gamma(n) = -d - e * (n-1)$
 - d: gap initiation penalty
 - e: gap extension penalty



Affine gap algorithm

- Dynamical programming in 3 layers
 - The three recurrences for the scoring algorithm creates a 3-layered graph.
 - The top level creates/extends gaps in the sequence w .
 - The bottom level creates/extends gaps in sequence v .
 - The middle level extends matches and mismatches.

- Keep track of 3 matrices



Affine Gap Update rule

$$\downarrow F_{i,j} = \max \begin{cases} \downarrow F_{i-1,j} - e & \text{Continue Gap in } s \text{ (deletion)} \\ F_{i-1,j} - (d+e) & \text{Start Gap in } s \text{ (deletion): from middle} \end{cases}$$

$$\vec{F}_{i,j} = \max \begin{cases} \vec{F}_{i,j-1} - e & \text{Continue Gap in } t \text{ (insertion)} \\ F_{i,j-1} - (d+e) & \text{Start Gap in } t \text{ (insertion):from middle} \end{cases}$$

$$F_{i,j} = \max \begin{cases} F_{i-1,j-1} + s(v_i, w_j) & \text{Match or Mismatch} \\ \downarrow F_{i,j} & \text{End deletion: from top} \\ \vec{F}_{i,j} & \text{End insertion: from bottom} \end{cases}$$

How to decide on the correct scoring metric

- Scoring metrics should reflect the evolutionary process
- What are the odds that an alignment is biologically meaningful – the proteins are homologous
- Random model: product of chance events
- Non-random model: two sequences derived from a common ancestor
- Things to consider
 - What is the frequency of different mutations
 - Over what time scale?

Log-odds scoring

What are the odds that this alignment is meaningful?

$$\begin{array}{c} X_1 X_2 X_3 \dots X_n \\ Y_1 Y_2 Y_3 \dots Y_n \end{array}$$

Random model: We're observing a chance event. The probability is

$$\prod_i p_{X_i} \prod_i p_{Y_i}$$

where p_X is the frequency of X

Alternative: The two sequences derive from a common ancestor. The probability is

$$\prod_i q_{X_i Y_i}$$

where q_{XY} is the joint probability that X and Y evolved from the same ancestor.

Log-odds scoring

Odds ratio:

$$\frac{\prod_i q_{X_i Y_i}}{\prod_i p_{X_i} \prod_i p_{Y_i}} = \prod_i \frac{q_{X_i Y_i}}{p_{X_i} p_{Y_i}}$$

Log-odds ratio (score):

$$S = \sum s(X_i, Y_i)$$

where

$$s(X, Y) = \log \left(\frac{q_{XY}}{p_X p_Y} \right)$$

is the score for X, Y . The $s(X, Y)$'s define a scoring matrix

Conservative substitutions

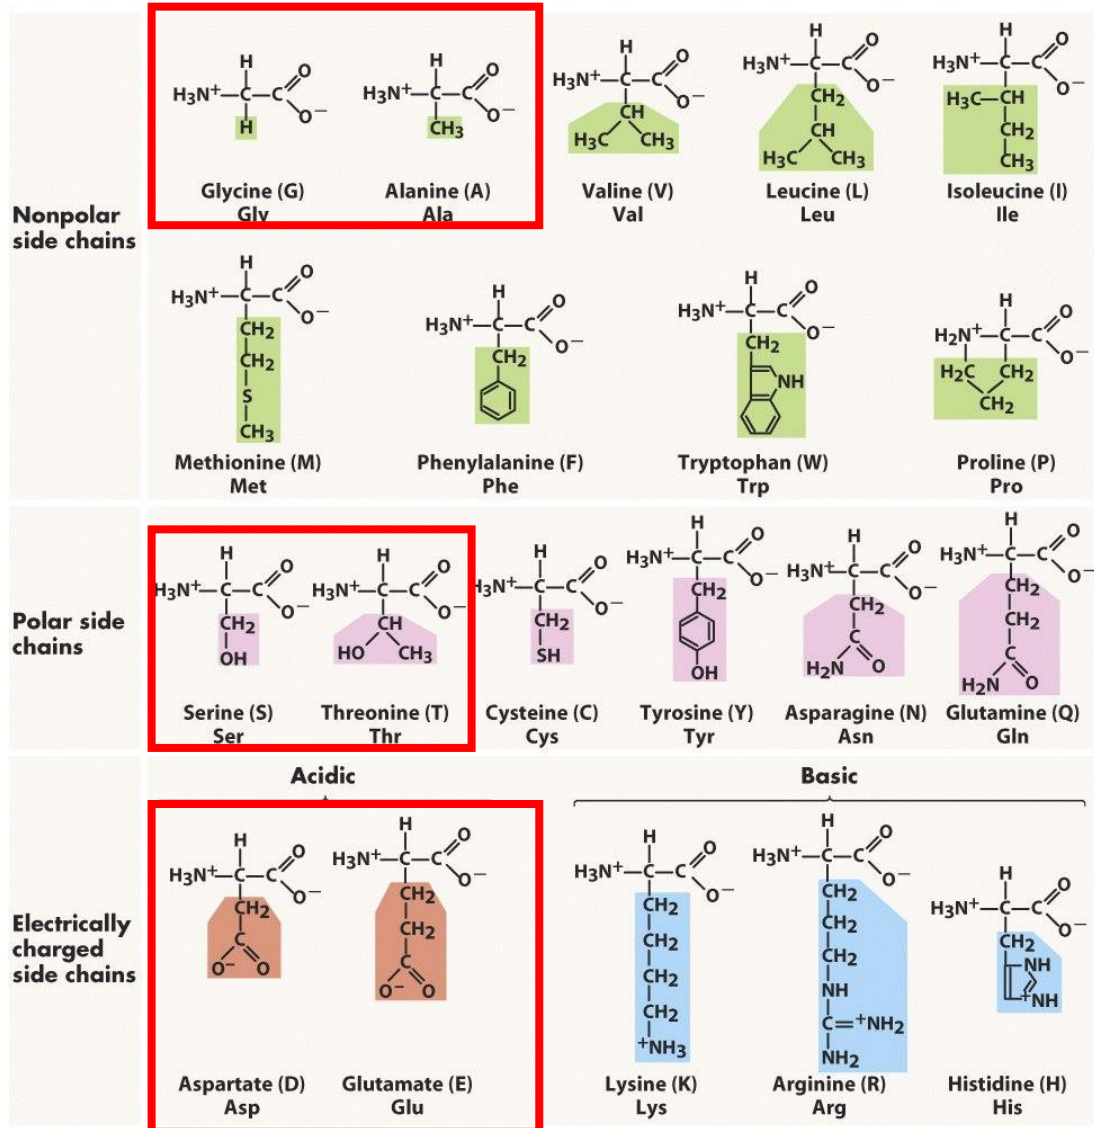


Figure 3-5 Biological Science, 2/e

PAM1 probability matrix

- PAM1 probability matrix
- Dayhoff et al (1978) estimated probability of one-step transitions
- Used a family of very closely related proteins
- Corresponds to 1 change per 100 a.a.

		Original amino acid																			
		A Ala	R Arg	N Asn	D Asp	C Cys	Q Gln	E Glu	G Gly	H His	I Ile	L Leu	K Lys	M Met	F Phe	P Pro	S Ser	T Thr	W Trp	Y Tyr	V Val
Replacement amino acid	A	98.7	0.0	0.1	0.1	0.0	0.1	0.2	0.2	0.0	0.1	0.0	0.0	0.1	0.0	0.2	0.4	0.3	0.0	0.0	0.2
	R	0.0	99.1	0.0	0.0	0.0	0.1	0.0	0.0	0.1	0.0	0.0	0.2	0.0	0.0	0.0	0.1	0.0	0.1	0.0	0.0
	N	0.0	0.0	98.2	0.4	0.0	0.0	0.1	0.1	0.2	0.0	0.0	0.1	0.0	0.0	0.0	0.2	0.1	0.0	0.0	0.0
	D	0.1	0.0	0.4	98.6	0.0	0.1	0.5	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0
	C	0.0	0.0	0.0	0.0	99.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0
	Q	0.0	0.1	0.0	0.1	0.0	98.8	0.3	0.0	0.2	0.0	0.0	0.1	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0
	E	0.1	0.0	0.1	0.6	0.0	0.4	98.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	G	0.2	0.0	0.1	0.1	0.0	0.0	0.1	99.4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.0	0.0	0.0	0.1
	H	0.0	0.1	0.2	0.0	0.0	0.2	0.0	0.0	99.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	I	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	98.7	0.1	0.0	0.2	0.1	0.0	0.0	0.1	0.0	0.0	0.3
	L	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.2	99.5	0.0	0.5	0.1	0.0	0.0	0.0	0.0	0.0	0.2
	K	0.0	0.4	0.3	0.1	0.0	0.1	0.1	0.0	0.0	0.0	0.0	99.3	0.2	0.0	0.0	0.1	0.1	0.0	0.0	0.0
	M	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.0	98.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	F	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.0	0.0	99.5	0.0	0.0	0.0	0.0	0.3	0.0
	P	0.1	0.1	0.0	0.0	0.0	0.1	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	99.3	0.1	0.0	0.0	0.0	0.0
	S	0.3	0.1	0.3	0.1	0.1	0.0	0.1	0.2	0.0	0.0	0.0	0.1	0.0	0.0	0.2	98.4	0.4	0.1	0.0	0.0
	T	0.2	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.1	0.1	0.0	0.1	0.3	98.7	0.0	0.0	0.1
	W	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	99.8	0.0	0.0
	Y	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.0	0.0	0.0	0.0	99.5	0.0
	V	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.6	0.1	0.0	0.2	0.0	0.0	0.0	0.1	0.0	0.0	99.0

FIGURE 3.9 The PAM1 mutation probability matrix. The original amino acid j is arranged in columns (across the top), while the replacement amino acid i is arranged in rows. Dayhoff et al. multiplied values by 10,000 (offering added precision) while here we multiply by 100 so that, for example, the first cell's value of 98.7 corresponds to 98.7% occurrence of ala remaining ala over this evolutionary interval.

PAM1 through PAM250

- We can multiply PAM1 by itself to get a probability matrix for longer time scales
- PAM is measured in number of changes not time
- Number of changes that occurred is not the same as number of observed changes

Observed differences in 100 residues	Evolutionary distance in PAMs
1	1.0
5	5.1
10	10.7
15	16.6
20	23.1
25	30.2
30	38.0
35	47
40	56
45	67
50	80
55	94
60	112
65	133
70	159
75	195
80	246

Source: Dayhoff (1972). Reproduced with permission from National Biomedical Research Foundation.


PAM250

- Only 20% identity
- 20% identity is close to what you might get aligning random sequences

		Original amino acid																			
		A	R	N	D	C	Q	E	G	H	I	L	K	M	F	P	S	T	W	Y	V
Replacement amino acid	A	13	6	9	9	5	8	9	12	6	8	6	7	7	4	11	11	11	2	4	9
	R	3	17	4	3	2	5	3	2	6	3	2	9	4	1	4	4	3	7	2	2
	N	4	4	6	7	2	5	6	4	6	3	2	5	3	2	4	5	4	2	3	3
	D	5	4	8	11	1	7	10	5	6	3	2	5	3	1	4	5	5	1	2	3
	C	2	1	1	1	52	1	1	2	2	2	1	1	1	1	2	3	2	1	4	2
	Q	3	5	5	6	1	10	7	3	7	2	3	5	3	1	4	3	3	1	2	3
	E	5	4	7	11	1	9	12	5	6	3	2	5	3	1	4	5	5	1	2	3
	G	12	5	10	10	4	7	9	27	5	5	4	6	5	3	8	11	9	2	3	7
	H	2	5	5	4	2	7	4	2	15	2	2	3	2	2	3	3	2	2	3	2
	I	3	2	2	2	2	2	2	2	2	10	6	2	6	5	2	3	4	1	3	9
	L	6	4	4	3	2	6	4	3	5	15	34	4	20	13	5	4	6	6	7	13
	K	6	18	10	8	2	10	8	5	8	5	4	24	9	2	6	8	8	4	3	5
	M	1	1	1	1	0	1	1	1	1	2	3	2	6	2	1	1	1	1	1	2
	F	2	1	2	1	1	1	1	1	3	5	6	1	4	32	1	2	2	4	20	3
	P	7	5	5	4	3	5	4	5	5	3	3	4	3	2	20	6	5	1	2	4
	S	9	6	8	7	7	6	7	9	6	5	4	7	5	3	9	10	9	4	4	6
	T	8	5	6	6	4	5	5	6	4	6	4	6	5	3	6	8	11	2	3	6
	W	0	2	0	0	0	0	0	0	1	0	1	0	0	1	0	1	0	55	1	0
	Y	1	1	2	1	3	1	1	1	3	2	2	1	2	15	1	2	2	3	31	2
	V	7	4	4	4	4	4	4	5	4	15	10	4	10	5	5	5	7	2	4	17


Choice of scale reflects the result

- Human and chimp beta globin—close orthologs
- Human beta and alpha globin – paralogs –further apart



PAM model

- Assumptions
 - Replacement at any site depends only on the a.a. on that site, give the **mutability** of the a.a.
 - Sequences in the training set (and those compared) have average a.a. composition.
- Sources of error
 - Many proteins depart from the average a.a. composition.
 - The a.a. composition can vary even within a protein (e.g. transmembrane proteins).
 - A.a. positions are not “mutated” equally probably; especially in lor evolutionary distances.
 - Rare replacements are observed too infrequently and...
 - ...errors in PAM1 are magnified in PAM250.




Blocks Substitution Matrices (BLOSUM):

- Log-likelihood matrix (Henikoff & Henikoff, 1992)
- BLOCKS database of aligned sequences used as primary source set.
- Different BOLSUM n matrices are calculated independently from BLOCKS (ungapped local alignments)
- BLOSUM n is based on a cluster of BLOCKS of sequences that share at least n percent identity
- BLOSUM62 represents closer sequences than BLOSUM45
- BLOCKS database contains large number of ungapped multiple local alignments of conserved regions of proteins
- Alignments include distantly related sequences in which multiple base substitutions at the same position could be observed

PAM vs BLOSUM

- PAM is based on closely related sequences, thus is biased for short evolutionary distances where number of mutations are scalable
- PAM is based on globally aligned sequences, thus includes conserved and non-conserved positions; BLOSUM is based on conserved positions only
- Lower PAM/higher BLOSUM matrices identify shorter local alignments of highly similar sequences
- Higher PAM/lower BLOSUM matrices identify longer local alignments of more distant sequences



- Matrices of choice:
 - BLOSUM62: the all-weather matrix
 - PAM250: for distant relatives