## Advanced Algorithms and Models for Computational Biology

## Modeling biological sequences

- Kinds of questions we want to ask
- How to align two sequence to reveal conserved regions?
- Is this sequence a motif (e.g., binding site, splice site)?
- is this sequence part of the coding region of a gene?
- Are these two sequences evolutionarily related?
$\bullet$
- What we will not address (covered last semester)
- How multiple sequences can be optimally aligned
- how sequencing results of a clone library can be assembled
- What is the most parsimonious phylogeny of a set of sequences
- Machine learning : extracting useful information from a corpus of data $D$ by building good (predictive, evaluative or decision) models


## Modeling biological sequences, ctd

- We will use probabilistic models of sequences -- not the only approach, but usually the most powerful, because
- sequences are the product of an evolutionary process which is stochastic in nature,
- want to detect biological "signal" against "random noise" of background mutations,
- data may be missing due to experimental reasons or intrinsically unobservable, and
- we want to integrate multiple (heterogeneous) data and incorporate prior knowledge in a flexible and principled way,
$\bullet$
- Computational analysis only generate hypothesis, which must be tested by experiments
- Site-directed mutagenesis (to alter the sequence content)
- Knockouts/insertions of genes/sites (deletion/addition of elements)
- Functional perturbations (pathway inhibitors, drugs, ...)
- From one-way learning to close-loop learning:
- Active learning: can a machine design smart experiments?


## Hierarchical structure of the genome



## The DNA strand has a chemical polarity



5' C-A-G 3'


## Writing DNA sequence

- One strand is written by listing its bases in $5^{\prime}$ to $3^{\prime}$ order

> 5' ACCGTTACT 3'

- Each strand uniquely determines the complementary strand, which runs in the opposite direction:

$$
\begin{aligned}
& 5^{\prime} \text { ACCGTTACT 3' } \\
& 3^{\prime} \text { ' }
\end{aligned}
$$

- So the reverse complement of ACCGTTACT is written TGGCAATGA
- In general people write one strand and in 5' to 3' order
- This is the ordering that a polymerase or a ribosome scan the sequence
- Establishes a common standard for genome nomenclatures



## Gene structure in prokaryotes

- A protein-coding gene consists of the following, in 5' to $3^{\prime}$ order
- An upstream regulatory region, generally < 50 bp , which turns transcription on and off.
- A transcription start site where RNA polymerase incorporates 1st nucleotide into nascent mRNA.
- A 5' untranslated region, generally < 30bp, that is transcribed into mRNA but not translated.
- The translation start site marking the start of the coding region. Consists of a start codon, which causes the start of translation
- The coding region of the gene (typically=1000bp), consisting of a sequence of codons.
- The translation stop site marking the end of coding region. Consists of a stop codon, which causes the release of the polypeptide at conclusion of translation.
- A 3' untranslated region, transcribed into RNA but not translated.
- The transcription stop site marking where the RNA polymerase concludes transcription.



## Gene structure in eukaryotes

- A typical gene consist of the following, in 5' to 3' order
- An upstream regulatory region, often larger and more complex than in prokaryotes, parts of which may be several thousand bases or more upstream of transcription start site.
- A transcription start site.
- A 5' untranslated region, often larger than in prokaryotes, and which may include sequences playing a role in translation regulation.
- The coding sequence, which unlike the case with prokaryotes, may be interrupted by non-coding regions called introns. These are spliced out of the transcript to form the mature mRNA (and sometimes the splicing can occur in more than one way).
- The translation stop site.
- A 3' untranslated region, which may contain sequences involved in translational regulation
- A polyadenylation (playA) signal, which indicates to the cell's RNA processing machinery that the RNA transcript is to be cleaved and a poly-adenine sequence (AAAAAA...) tail appended to it
- The transcription stop site.

mRNA1

|  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- |

mRNA2


## Eukaryotic genome structure

- Genes may be transcribed in either direction, and can overlap

gene 2
gene 4



## Sequence alignment:





## Sequence conservation implies functional conservation



Alignment is the key to

- Finding important regions
- Determining function
- Uncovering the evolutionary forces


## Sequence-based functional prediction

- Sequence similarity is useful in predicting the function of a new sequence...
- ... assuming that sequence similarity implies structural and functional similarity.



## Sequence Alignment

AGGCTATCACCTGACCTCCAGGCCGATGCCC
TAGCTATCACGACCGCGGTCGATTTGCCCGAC


- AGGCTATCACCTGACCTCCAGGCCGA - - TGCCC - -

TAG-CTATCAC--GACCGC--GGTCGATTTGCCCGAC

## Definition

Given two strings

$$
x=x_{1} x_{2} \ldots x_{M}, \quad y=y_{1} y_{2} \ldots y_{N},
$$

An alignment of two sequences $\boldsymbol{x}$ and $\boldsymbol{y}$ is an arrangement of $\boldsymbol{x}$ and $\boldsymbol{y}$ by position, where $\boldsymbol{a}$ and $\boldsymbol{b}$ can be padded with gap symbols to achieve the same length.

## Editing Distance

- Sequence edits:
- Mutations

AGGCCTC

- Insertions AGGGCCTC
- Deletions

AGG_CTC

- We can turn the edit protocol into a measure of distance by assigning a "cost" or "weight" $S$ to each operation.
- For example, for arbitrary characters $u, v$ from set $A$ we may define
$S(u, u)=0 ; S(u, v)=1$ for $u \neq v ; S(u,-)=S(-, v)=1$. (Unit Cost)
- This scheme is known as the Levenshtein distance, also called unit cost model. Its predominant virtue is its simplicity.
- In general, more sophisticated cost models must be used.
- For example, replacing an amino acid by a biochemically similar one should weight less than a replacement by an amino acid with totally different properties.


## Scoring Function

- Scoring Function:

| Match: | +m |
| :--- | :--- |
| Mismatch: | $-\mathrm{s} \quad$ (a more sophisticated score matrix can be used for proteins) |
| Gap: | -d |

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

- The Alignment Score of $\boldsymbol{x}$ and $\boldsymbol{y}$ is the score of an optimal alignment of $\boldsymbol{x}$ and $\boldsymbol{y}$ under a score function $S$. We denote it by $F(x, y)$.
- For example, using the score function corresponding to the unit cost model in our previous example, we obtain the following score:

```
a: AGCACAC-A or AG-CACACA
b: A-CACACTA ACACACT-A
cost: -2
    cost: -4
```

- Here it is easily seen that the left-hand assignment is optimal under the unit cost model, and hence the alignment score $F(\mathbf{a}, \boldsymbol{b})=-2$.


## Scoring Matrices

- Physical/Chemical similarities
- comparing two sequences according to the properties of their residues may highlight regions of structural similarity
- The matrix that performs best will be the one that best reflects the evolutionary separation of the sequences being aligned
- The most commonly used mutation matrices: PAM or BLOSUM

Below diagonal: BLOSUM62 substitution matrix
Above diagonal: Difference matrix obtained by subracting the PAM 160 matrix entrywise.
(Henikoff \& Henikoff 1992)


## How do we compute the best alignment?

- A alignment corresponds to a path in the alignment matrix



Too many possible alignments:

$$
\mathrm{O}\left(2^{\mathrm{M}+\mathrm{N}}\right)
$$

## Dynamic Programming

- The optimum alignment is obtained by tracing the highest scoring path from the top left-hand corner to the bottom righthand corner of the matrix (or the lowest editing-distance path from bottom right-hand corner to top left-hand corner)
- When the alignment steps away from the diagonal this implies an insertion or deletion event, the impact of which can be assessed by the application of a gap penalty
- Dynamic Programming: recursively solve nested problems each of a manageable size


## Dynamic Programming

- Three possible cases:

1. $x_{i}$ aligns to $y_{j}$

$$
x_{1} \ldots \ldots x_{i-1} \quad x_{i}
$$

$$
y_{1 \ldots \ldots} \ldots y_{j-1} \quad y_{j}
$$

$$
F(i, j)=F(i-1, j-1)+\left\{\begin{array}{l}
m, \text { if } x_{i}=y_{j} \\
-s, \text { if not }
\end{array}\right.
$$

2. $x_{i}$ aligns to a gap

$$
\begin{aligned}
& x_{1} \ldots \ldots x_{i-1} x_{i} \\
& y_{1} \ldots \ldots y_{j}-
\end{aligned} \quad F(i, j)=F(i-1, j)-d
$$

3. $y_{j}$ aligns to a gap

$$
\begin{array}{ll}
x_{1} \ldots \ldots x_{i}- & F(i, j)=F(i, j-1)-d \\
y_{1} \ldots \ldots y_{i-1} & y_{i}
\end{array}
$$

## Dynamic Programming (cont'd)

- How do we know which case is correct?

Inductive assumption:

$$
F(i, j-1), F(i-1, j), F(i-1, j-1) \quad \text { are optimal }
$$

Then,

$$
F(i, j)=\max \left\{\begin{array}{l}
F(i-1, j-1)+s\left(x_{i,} y_{j}\right) \\
F(i-1, j)-d \\
F(i, j-1)-d
\end{array}\right.
$$

Where

$$
s\left(x_{i}, y_{j}\right)=m \text {, if } x_{i}=y_{j} ;-s, \text { if not }
$$

## Example

$$
\begin{aligned}
& x=\text { AGTA } \\
& y=\text { ATA }
\end{aligned}
$$

$$
m=1
$$

$$
s=-1
$$

$$
d=-1
$$

$$
F(i, j) \quad i=0 \quad 1 \quad 2 \quad 3 \quad 4
$$

|  |  |  | $A$ | $G$ | $T$ | $A$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  |  |  |  |  |  |

Optimal Alignment:

$$
F(4,3)=2
$$

AGTA
A - TA

## Alignment is additive

- Observation:

The score of aligning

$$
\begin{aligned}
& x_{1} \ldots \ldots x_{M} \\
& y_{1} \ldots \ldots y_{N}
\end{aligned}
$$

is additive

| Say that | $x_{1 \ldots} \ldots x_{i}$ | $x_{i+1} \ldots x_{M}$ |
| :--- | :--- | :--- |
| aligns to | $y_{1 \ldots y_{j}}$ | $y_{j+1 \cdots} y_{N}$ |

The two scores add up:

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

$$
F^{*}\left(x[1: M], y[1: M)=\operatorname{Max}_{i j}\left\{F^{*}(x[1:], y[1: j])+F^{*}(x[i+1: M], y[j+1: M)\}\right.\right.
$$



## The Needleman-Wunsch Algorithm

1. Initialization.
a. $F(0,0)$
$=0$
b. $F(0, j)$
$=-j \times d$
c. $F(i, 0)$
$=-\mathrm{i} \times \mathrm{d}$
2. Main Iteration. Filling-in partial alignments
a. For each $\mathrm{i}=1 \ldots . . \mathrm{M}$

For each $\quad j=1 \ldots \ldots . N$

$$
\begin{aligned}
& \mathrm{F}(\mathrm{i}, \mathrm{j})=\max \begin{cases}\mathrm{F}(\mathrm{i}-1, \mathrm{j}-1)+\mathrm{s}\left(\mathrm{x}_{\mathrm{i}}, \mathrm{y}_{\mathrm{j}}\right) & \text { [case 1] } \\
\mathrm{F}(\mathrm{i}-1, \mathrm{j})-\mathrm{d} & \text { [case 2] } \\
\mathrm{F}(\mathrm{i}, \mathrm{j}-1)-\mathrm{d} & \text { [case 3] }\end{cases} \\
& \mathrm{Ptr}(\mathrm{i}, \mathrm{j})= \text { DIAG, }, \begin{array}{l}
\text { if [case 1] }
\end{array} \\
& \text { LEFT, } \begin{array}{l}
\text { if [case 2] } \\
\text { if [case 3] }
\end{array}
\end{aligned}
$$

3. Termination. $\mathrm{F}(\mathrm{M}, \mathrm{N})$ is the optimal score, and from $\operatorname{Ptr}(\mathrm{M}, \mathrm{N})$ 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:
----------CTATCACCTGACCTCCAGGCCGATGCCCCTTCCGGC GCGAGTTCATCTATCAC--GACCGC--GGTCG-------------
- Then, we don't want to penalize gaps in the ends
- Different types of overlaps



## The Overlap Detection variant



Changes:

1. Initialization

For all $\mathrm{i}, \mathrm{j}$,
$F(i, 0)=0$
$F(0, j)=0$
2. Termination
$F_{\text {OPT }}=\max \left\{\begin{array}{l}\max _{\mathrm{i}} \mathrm{F}(\mathrm{i}, \mathrm{N}) \\ \max _{\mathrm{j}} \mathrm{F}(\mathrm{M}, \mathrm{j})\end{array}\right.$

## The local alignment problem

- The problem:
- Given two strings $x=x_{1} \ldots \ldots x_{M}$,
$y=y_{1} \ldots . . y_{N}$
- Find substrings $x^{\prime}, y^{\prime}$ whose similarity (optimal global alignment value) is maximum
- e.g. $x=$ aaaacccccgggg

$$
y=\text { cccgggaaccaacc }
$$

- Why
- Genes are shuffled between genomes
- Portions of proteins (domains) are often conserved



## The Smith-Waterman algorithm

Idea: Ignore badly aligning regions

Modifications to Needleman-Wunsch:


Iteration:
$F(i, j)=\max \left\{\begin{array}{l}0 \\ F(i-1, j)-d \\ F(i, j-1)-d \\ F(i-1, j-1)+s\left(x_{i}, y_{j}\right)\end{array}\right.$

## The Smith-Waterman algorithm

## Termination:

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

$$
F_{\mathrm{OPT}}=\max _{\mathrm{i}, \mathrm{j}} \mathrm{~F}(\mathrm{i}, \mathrm{j})
$$

2. If we want all local alignments scoring >t
?? $\quad$ For all $\mathrm{i}, \mathrm{j}$ find $\mathrm{F}(\mathrm{i}, \mathrm{j})>\mathrm{t}$, and trace back

Complicated by overlapping local alignments

## Scoring the gaps more accurately

- Current model:
- Gap of length $n$
- incurs penalty $n \times d$
$\gamma(\mathrm{n})$

- However, gaps usually occur in bunches
- Convex (saturating) gap penalty function:
$\gamma(\mathrm{n})$ :
$\gamma(\mathrm{n})$

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


## Convex gap dynamic programming

Iteration:
$F(i, j)=\max \left\{\begin{array}{l}F(i-1, j-1)+s\left(x_{i}, y_{j}\right) \\ \max _{k=0 \ldots i-1} F(k, j)-\gamma(i-k) \\ \max _{k=0 . . j-1} F(i, k)-\gamma(j-k)\end{array}\right.$

Termination: same

Running Time: $\mathrm{O}\left(\mathrm{N}^{2} \mathrm{M}\right)$
(assume $\mathrm{N}>\mathrm{M}$ )
Space: O(NM)

## Compromise: affine gaps

- Simple piece-wise linear gap penalty

| $\gamma(\mathrm{n})=\mathrm{d}+(\mathrm{n}-1) \times \mathrm{e}$ |  |
| :---: | :--- |
| $\quad \mid$ | $\mid$ |
| gap | gap |
| open | extend |



- Fancier Piece-wise linear gap penalty

- Think of how you would compute optimal alignment with this gap function in $\mathrm{O}(\mathrm{MN})$


## Bounded Dynamic Programming

- Assume we know that $x$ and $y$ are very similar

Assumption: \# gaps $(\mathrm{x}, \mathrm{y})<\mathrm{k}(\mathrm{N}) \quad$ ( say $\mathrm{N}>\mathrm{M}$ )

Then, $\quad \begin{aligned} & x_{i} \\ & y_{j}\end{aligned} \quad$ implies $\quad|i-j|<k(N)$

We can align $x$ and $y$ more efficiently:

Time, Space:

$$
\mathrm{O}(\mathrm{~N} \times \mathrm{k}(\mathrm{~N})) \ll \mathrm{O}\left(\mathrm{~N}^{2}\right)
$$



## State of biological databases



## State of biological databases

- Number of genes in these genomes:
- Mammals: $\quad \sim 25,000$
- Insects: $\sim 14,000$
- Worms: ~17,000
- Fungi: $\sim 6,000-10,000$
- Small organisms: 100s-1,000s
- Each known or predicted gene has one or more associated protein sequences
- >1,000,000 known / predicted protein sequences


## Some useful applications of alignments

- Given a newly discovered gene,
- Does it occur in other species?
- How fast does it evolve?
- Assume we try Smith-Waterman:


The entire genomic database

$$
10^{10}-10^{12}
$$

## Some useful applications of alignments

- Given a newly sequenced organism,
- Which subregions align with other organisms?
- Potential genes
- Other biological characteristics
- Assume we try Smith-Waterman:



## Indexing-based local alignment

- BLAST- Basic Local Alignment Search Tool

Running Time: O(MN)
However, orders of magnitude faster than SmithWaterman

## Main idea:

1. Construct a dictionary of all the words in the query
2. Initiate a local alignment for each word match between query and DB

## Multiple alignment

- The simultaneous alignment of a number of DNA or protein sequences is one of the commonest tasks in bioinformatics.
- Useful for:
- phylogenetic analysis (inferring a tree, estimating rates of substitution, etc.)
- detection of homology between a newly sequenced gene and an existing gene family
- prediction of protein structure
- demonstration of homology in multigene families
- determination of a consensus sequence (e.g., in assembly)
- Can we naively use DP?
- need to deal with k-dimensional table for k sequences ...


## Extending the pairwise alignment algorithms

- Generally not feasible for more than a small number of sequences ( $\sim 5$ ), as the necessary computer time and space quickly becomes prohibitive.
- Computational time grows as $\mathrm{N}^{m}$, where $m=$ number of sequences.
- For example, for 100 residues from 5 species, $1005=10,000,000,000$ (i.e., the equivalent of two sequences each 100,000 residues in length.)
- Nor is it wholly desirable to reduce multiple alignment to a similar mathematical problem to that tackled by pairwise alignment algorithms.
- Two issues which are important in discussions of multiple alignment are:
- the treatment of gaps: position-specific and/or residue-specific gap penalties are both desirable and feasible, and
- the phylogenetic relationship between the sequences (which must exist if they are alignable): it should be exploited.


## Progressive alignment

- Up until about 1987, multiple alignments would typically be constructed manually, although a few computer methods did exist.
- Around that time, algorithms based on the idea of progressive alignment appeared.
- In this approach, a pairwise alignment algorithm is used iteratively,
- first to align the most closely related pair of sequences,
- then the next most similar one to that pair, and so on.
- The rule "once a gap, always a gap" was implemented, on the grounds that the positions and lengths of gaps introduced between more similar pairs of sequences should not be affected by more distantly related ones.
- The most widely used progressive alignment algorithm is currently CLUSTAL W.
- Other methods include the profile HMM-based methods

- The three basic steps in the CLUSTAL W approach are shared by all progressive alignment algorithms:
A. Calculate a matrix of pairwise distances based on pairwise alignments between the sequences
B. Use the result of A to build a guide tree, which is an inferred phylogeny for the sequences
C. Use the tree from B to guide the progressive alignment of the sequences
- We will omit details


## Web-based multiple sequence alignment

- ClustalW
- www2.ebi.ac.uk/clustalw/
- dot.imgen.bcm.tmc.edu:9331/multi-align/Options/clustalw.html
- www.clustalw.genome.ad.jp/
- bioweb.pasteur.frlintro-uk.html
- pbil.ibcp.fr
- transfac.gbf.de/programs.html
- www.bionavigator.com
- PileUp
- helix.nih.gov/newhelix
- www.hgmp.mrc.ac.uk/
- bcf.arl.arizona.edu/gcg.html
- www.bionavigator.com
- Dialign
- genomatix.gsf.del
- bibiserv.techfak.uni-bielefeld.del
- bioweb.pasteur.frlintro-uk.html
- www.hgmp.mrc.ac.uk/
- Match-box
- www.fundp.ac.be/sciences/biologie/bms/matchbox_submit.html
- For reviews: G. J. Gaskell, BioTechniques 2000, 29:60, and
- www.techfak.uni-bielefeld.de/bcd/Curric/MulAli/welcome.html


## Acknowledgments

- Serafim Batzoglou: for many of the slides adapted or modified from his lecture slides at Stanford University
- Terry Speed: for some of the slides modified from his lectures at UC Berkeley
- Phil Green and Joe Felsenstein: for some of the slides modified from his lectures at Univ. of Washington

