

Alignment and blast: basic principles and limitations to keep in mind for downstream analyses

David Dylus

ESCMID Postgraduate technical workshop

10/09/2019

Outline

- Motivation
 - Simple alignment
 - Complexity
- Querying a sequence against a database (BLAST)
 - Local vs Global alignment
 - How does it work?
 - What does the output mean?
- Mapping a sequence against a genome
 - Burrows Wheeler Transform
 - Tool Performance
 - Things to consider
- Multiple Sequence alignment
 - Basic Idea
 - Which tool is good for which purpose
- Conclusion

Why do you need alignment?

- Determine the expression of genes
 - Count the number of genes mapped to specific regions in the genome
- Find SNPs/SNVs in genomes
- Determine homology of genes
- Find the distance between two sequences
- Multiple sequence alignment for phylogenetic analysis

Simple alignment

- Where is GATTACA (len = 7)?

T G A T T A C A G A T T A C C (len = 15)

- 1.) Align at the beginning
- 2.) Compare each nucleotide with each nucleotide and count the number of matches
- 3.) Move one position along the genome and go back to 2.)

Simple alignment

- Where is GATTACA?

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	...
T	G	A	T	T	A	C	A	G	A	T	T	A	C	C	...
G	A	T	T	A	C	A									

Match Score: 1/7

First iteration = 7 comparisons
Total = 7

Simple alignment

- Where is GATTACA?

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	...
T	G	A	T	T	A	C	A	G	A	T	T	A	C	C	...
	G	A	T	T	A	C	A								

Match Score: 7/7

Second iteration = 7 comparisons
Total = 14

Simple alignment

- Where is GATTACA?


1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	...
T	G	A	T	T	A	C	A	G	A	T	T	A	C	C	...
		G	A	T	T	A	C	A	...						

Match Score: 1/7

Third iteration = 7 comparisons
Total = 21

Simple alignment

- Where is GATTACA?



1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	...
T	G	A	T	T	A	C	A	G	A	T	T	A	C	C	...
								G	A	T	T	A	C	A	

Match Score: 6/7 <- We may be very interested in these imperfect matches
Especially if there are no perfect end-to-end matches

9th iteration = 7 comparisons
Total = 63

$$\text{num_queries} * (\text{len_query} * (\text{len_genome} - \text{len_query} + 1))$$

What do we see?

- Brute force will take a long time for many queries
- Indels will be characterized as mismatches and alignment might be scored wrongly
- If we want to find alignments with small number of mismatches we need to store all of these, therefore not only runtime problem but also memory

Alignment

Local Alignment

Target Sequence

5' ACTACTAGATTACTTACGGATCAGGTACTTTAGAGGCTTGCAACCA 3'

||||| ||||| ||||| ||||| |||||

Query Sequence

5' TACTCACGGATGAGGTACTTTAGAGGC 3'

Global Alignment

Target Sequence

5' ACTACTAGATTACTTACGGATCAGGTACTTTAGAGGCTTGCAACCA 3'

||||| ||||| ||||| |||||

5' ACTACTAGATT---ACGGATC--GTACTTTAGAGGCTAGCAACCA 3'

Query Sequence

Alignment

Global Sequence Alignment	Local Sequence Alignment
In global alignment, an attempt is made to align the entire sequence (end to end alignment)	Finds local regions with the highest level of similarity between the two sequences.
A global alignment contains all letters from both the query and target sequences	A local alignment aligns a substring of the query sequence to a substring of the target sequence.
If two sequences have approximately the same length and are quite similar, they are suitable for global alignment.	Any two sequences can be locally aligned as local alignment finds stretches of sequences with high level of matches without considering the alignment of rest of the sequence regions.
Suitable for aligning two closely related sequences.	Suitable for aligning more divergent sequences or distantly related sequences.
Global alignments are usually done for comparing homologous genes like comparing two genes with same function (in human vs. mouse) or comparing two proteins with similar function.	Used for finding out conserved patterns in DNA sequences or conserved domains or motifs in two proteins.
A general global alignment technique is the Needleman–Wunsch algorithm.	A general local alignment method is Smith–Waterman algorithm.
Examples of Global alignment tools.	Examples of Local alignment tools.
<ul style="list-style-type: none"> • EMBOSS Needle • Needleman-Wunsch Global Align Nucleotide Sequences (Specialized BLAST) 	<ul style="list-style-type: none"> • BLAST • EMBOSS Water • LALIGN

BLAST

Build database with all 3 letter words part of query

N K C K T **P Q G** Q R L V N Q W N K

20 Amino acids
3 letter words

Word	Position
NKC	1
KCK	2
CKT	3
...	
PQG	6
...	

BLAST

Once a HSP is found extend the seed using dynamic programming until the score drops below a specific threshold and report the output


```
Query: N K C K T P Q G Q R L V N Q W N K
DB:    D S C V T P E G S R M L K R W D S

      2 -4 10 -9 7 8 1 6 -5 8 1 -2 -1 -2 13 2 -4
```

Score	Expect	Identities	Positives	Gaps
18.9 bits(37)	1e-04	6/14(43%)	9/14(64%)	0/14(0%)

Query	3	CKTPQGQRLVNQWN	16
		C TP+G R+ W+	
Sbjct	3	CVTPEGSRMLKRWD	16

BLAST



BLAST Stats

- Because we are using a matrix with different values for the alignment between two residues score and length of the alignment are not comparable
- E-value depends on database size and especially when using custom database can be misleading

Score	Expect	Identities	Positives	Gaps
18.9 bits(37)	1e-04	6/14(43%)	9/14(64%)	0/14(0%)
Query	3	CKTPQGQRLVNQWN	16	
		C TP+G R+ W+		
Sbjct	3	CVTPEGSRMLKRWD	16	

BLAST

Program	Query Sequence	Target Sequence
BLASTN	Nucleotide	Nucleotide
BLASTP	Protein	Protein
BLASTX	Nucleotide, six-frame translation	Protein
TBLASTN	Protein	Nucleotide, six-frame translation
TBLASTX	Nucleotide, six-frame translation	Nucleotide, six-frame translation

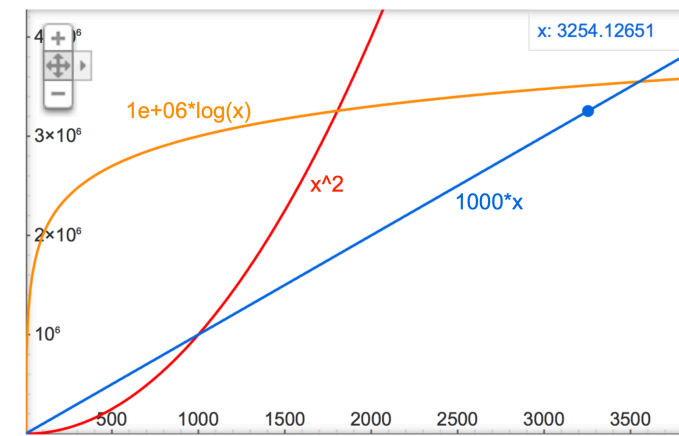
BLAST - Summary

- Searches for high-scoring segment pairs (HSPs)
 - Look for high scoring words of length W
 - Compile list L of all W -mers that score $>T$ with some word in query sequence
 - Scan database for words in L
 - When some word found: Extend alignment
 - When score drops more than X below hitherto best score stop extension
 - Report all words with large score S
- Results in all plausible alignments between two sequences

Questions BLAST

- Does BLAST give you the best possible alignment?
- Is the e-value a good choice to measure how well aligned it is?
- Would you use blast to align reads to a genome?

BLAST is great but SLOW



$O(\text{Number Queries} * \text{Length Query} * \text{DB Size})$

Typical:

Align 1.000.000.000 reads with 100bp length to
genome with 3.000.000 bp (Bacteria)

Alignment Problem: reads against Genome

- Detection of SNV
- Detection of SNPs
- Detection of presence
- Expression of genes / regions

How to align sequence reads to a reference?

Seeding: for each position, find longest exact match covering the position

Index genome with Burrows-Wheeler Transform

Align reads against genome

Burrows Wheel Transform (BWT):

- Computational trick on how data can be stored and searched
- Write all permutations of string and sort alphabetically (\$ first)



Current Pos	Nucleotide	Old Pos
1	g ₁	7
2	c ₁	5
3	\$	1
4	a ₁	2
5	a ₂	3
6	a ₃	4
7	c ₂	6

Align reads against genome

Burrows Wheel Transform (BWT):

- From last column alone we are able to reconstruct the whole genome

Current Pos	Nucleotide	Old Pos
1	g ₁	7
2	c ₁	5
3	\$	1
4	a ₁	2
5	a ₂	3
6	a ₃	4
7	c ₂	6


(b)



Align reads against genome

Burrows Wheel Transform (BWT):

- Exact matching can be performed in $O(\text{len query seq})$ time
- We know that a [2:4], c [5:6], g [7]



Current Pos	Nucleotide	Old Pos
1	g ₁	7
2	c ₁	5
3	\$	1
4	a ₁	2
5	a ₂	3
6	a ₃	4
7	c ₂	6

How to align sequence reads to a reference?

Extend seed (allow both local and end-to-end alignments with inexact matching)



Famous mappers

BWA (Li and Durbin 2009)

Bowtie2 (Langemead et al. 2009)

Slide kindly provided by Aitana Lebrand

Align reads against genome

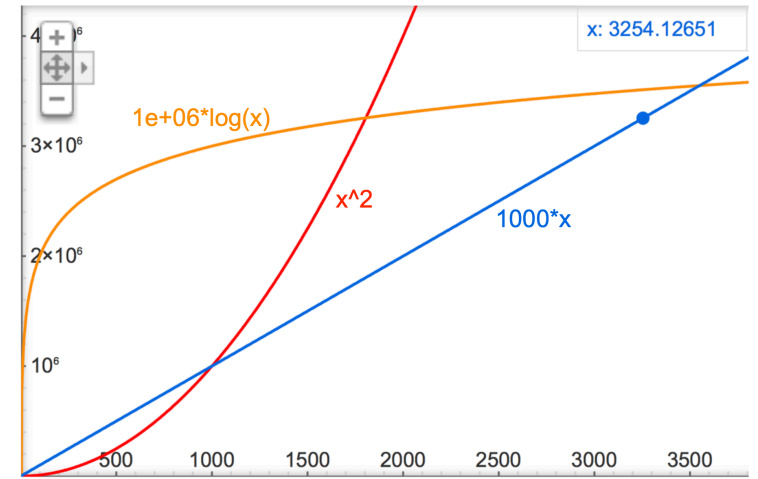
Burrows Wheel Transform (BWT):

- Inexact matching uses extend trick where at positions that not overlap quality of base call is considered and possible substitutions

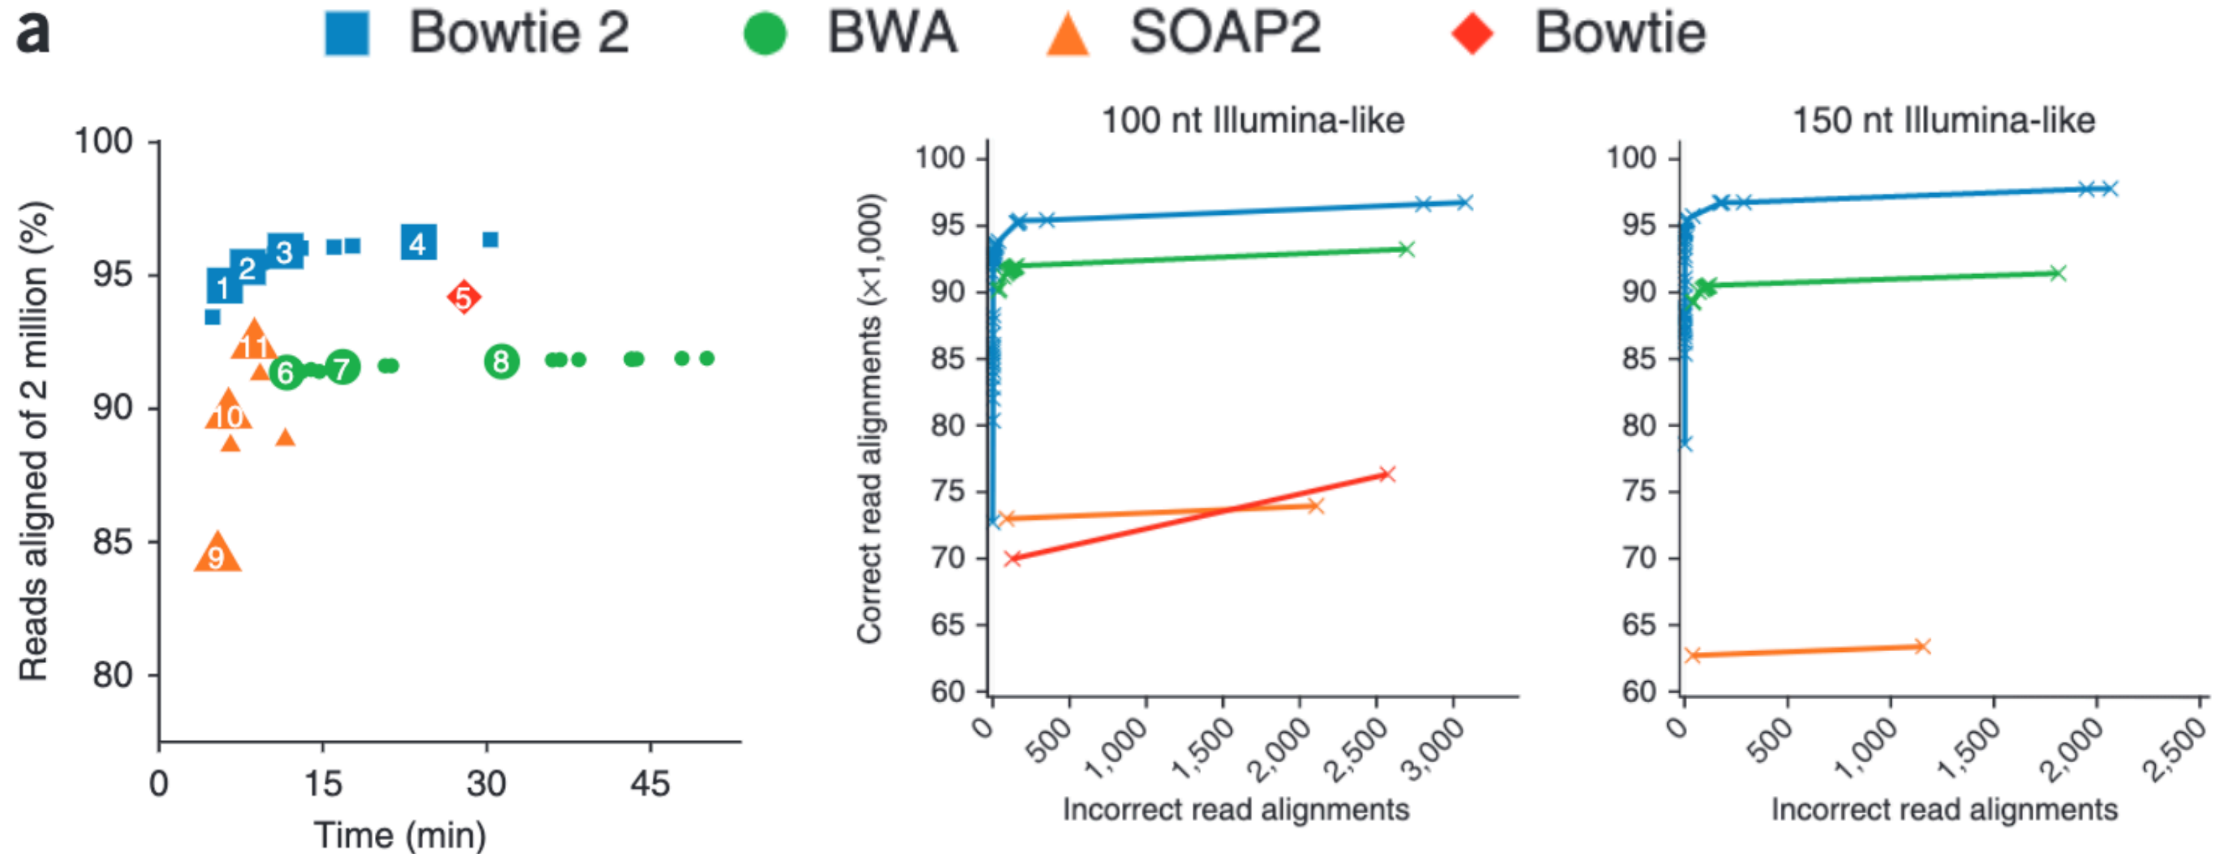
Align reads against genome

Burrows Wheel Transform (BWT):

- Transformation possible in $O(n)$
- Reconstruct Genome from BWT(Genome) in time $O(|\text{genome}|)$
- Search for all exact occurrences of read in time $O(|\text{read}|)$
- BWT(Genome) is easier to compress than Genome



Comparison of Tools



Be aware that this benchmark was done for Bowtie2 and other benchmarks might be different

Bowtie 2 typical run

```
[ddylus@dbc-serv05 bowtie2]$ time bowtie2 -x ref_genome -1 ../01_R1.fastq.gz -2 ../01_R2
.fastq.gz -S ref_genome.sam -p 6 --no-unal
945064 reads; of these:
  945064 (100.00%) were paired; of these:
    238691 (25.26%) aligned concordantly 0 times
    685378 (72.52%) aligned concordantly exactly 1 time
    20995 (2.22%) aligned concordantly >1 times
    ----
    238691 pairs aligned concordantly 0 times; of these:
      40883 (17.13%) aligned discordantly 1 time
    ----
    197808 pairs aligned 0 times concordantly or discordantly; of these:
      395616 mates make up the pairs; of these:
        353801 (89.43%) aligned 0 times
        36824 (9.31%) aligned exactly 1 time
        4991 (1.26%) aligned >1 times
81.28% overall alignment rate

real    0m35.467s
user    3m29.804s
sys     0m11.765s
```

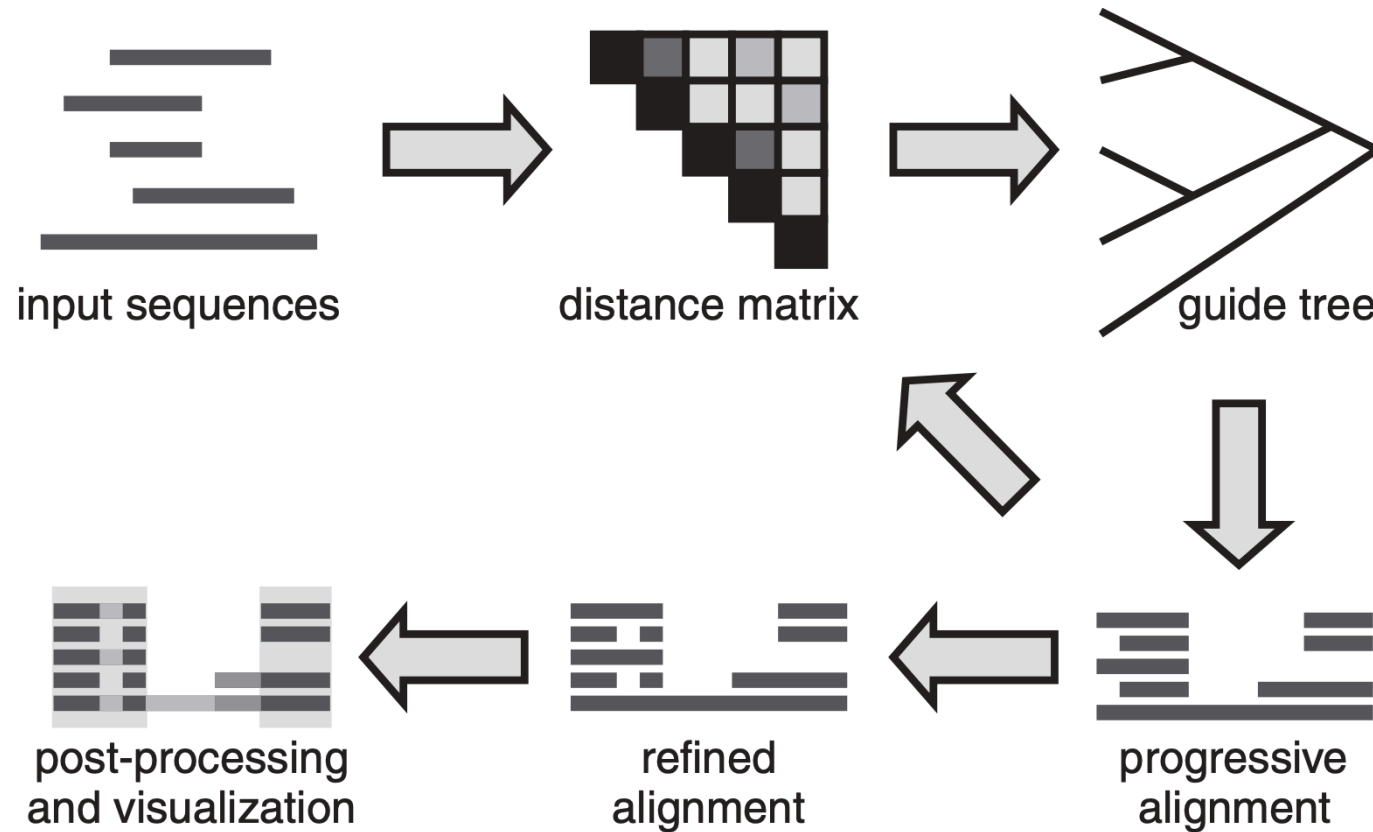
BWA typical runs

```
[M::mem_process_seqs] Processed 148366 reads in 11.650 CPU sec, 0.976 real sec  
[main] Version: 0.7.17-r1188  
[main] CMD: bwa mem -t 12 ../reference_genome.fasta ../01_R1.fastq.gz ../01_R2.fastq.gz  
[main] Real time: 15.105 sec; CPU: 136.568 sec  
  
real    0m15.165s  
user    2m13.188s  
sys     0m3.435s
```

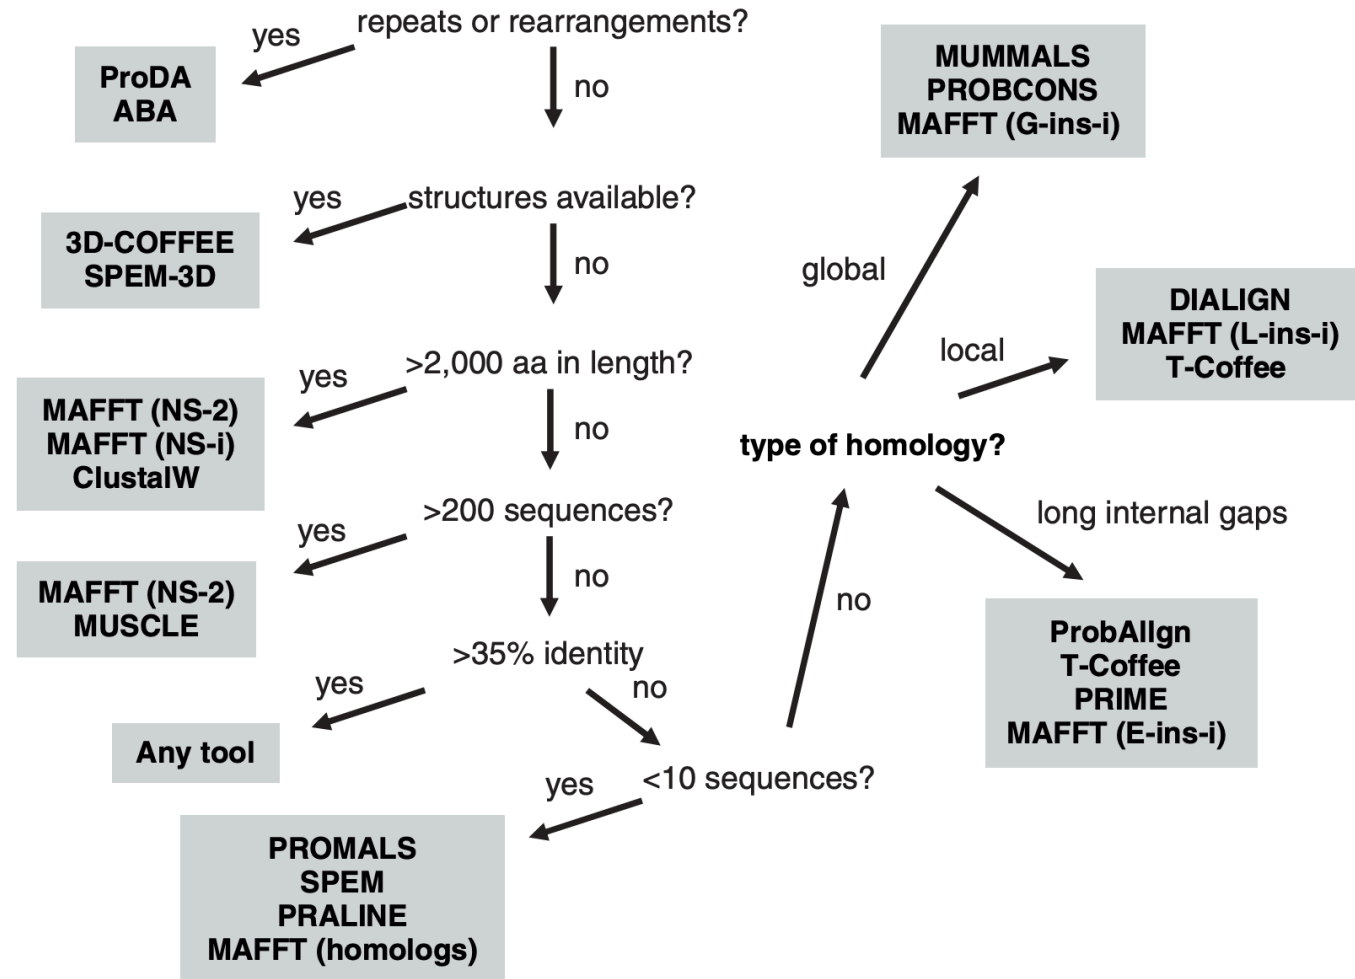
Aligning the whole read... or part of it

- Real life datasets are often not perfect:
- poor base call qualities, sequencing errors, insertions or deletions, structural variation, contaminants or adapter...
- **Despite adapter/quality trimming and allowing for mismatches/indels, some parts of the reads may still consist of sequences not present in the genome**
- **Hard-clipping** (removed from BAM) **vs. soft-clipping** (still part of BAM; not used for SVN calling, depth calculation, but can be useful to look into translocations, deletions...)
- **! Neither soft nor hard clipped regions are displayed in a viewer!**

What about Multiple Sequence Alignment



What about Multiple Sequence Alignment



Take home messages

- Be careful when comparing a blast result when using different search DBs
- Blast does not guarantee the most optimal alignment between you query and the obtained sequence from the DB
- Be aware that depending on read alignment tool you might end up with differences (for instance different SNPs)
- Do not just trust the MSA algorithm that is readily presented to you but decide based on your data

Source for this talk

Computational Biology:

- <https://ocw.mit.edu/courses/biology/7-91j-foundations-of-computational-and-systems-biology-spring-2014/video-lectures/>

Bowtie and BWA:

- <http://merenlab.org/2015/06/23/comparing-different-mapping-software/>

BWT:

- <https://www.youtube.com/watch?v=4n7NPk5lwbl>

Alignment:

- <https://www.youtube.com/watch?v=hpb-mH-yjLc&list=PL2mpR0RYFQsBiCWVJSvVAO3OJ2t7DzoHA>

Blast:

- <https://www.ndsu.edu/pubweb/~mcclean/plsc411/Blast-explanation-lecture-and-overhead.pdf>
- <https://developer.ibm.com/articles/j-seqalign/>
- <http://csc.columbusstate.edu/carroll/7840/private/papers/BasicLocalAlignmentSearchTool-BLAST.pdf>
- <https://www.youtube.com/watch?v=SAweFv8l8ow>
- http://web.math.ku.dk/~richard/courses/binf_project/Stinus-BLAST.pdf

References

MSA Benchmark

<https://arxiv.org/pdf/1211.2160.pdf>

<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2995051/>

<https://academic.oup.com/bioinformatics/advance-article/doi/10.1093/bioinformatics/btz552/5530966>

Mapper benchmark:

<https://www.ecseq.com/support/benchmark.html>

<https://www.biostars.org/p/125020/>