


Short read alignment (using external tools)

Kasper Daniel Hansen <khansen@jhsph.edu>
Brixen, 26 June 2011

Many slides are courtesy of
Hector Corrada Bravo and Ben Langmead

Analyzing reads




TATGTCGCAGTATCTG
TATGTCGCAGTATCTG
CCGGACACCCTATAT GTCGCAGTATCTGTCT
ACACCCTATGTCGCA
TATGTCGCAGTATCTT GTCGCAGTATCTGTNN
ACACCCTATGTCGCA
CCGGACACCCTATAT TATGTCGCAGTATCTG
GTCGCAGTATCTGTNN
TATGTCGCAGTATCTG
CCGGACACCCTATAT GTCGCAGTATCTGTC
GTCGCAGTATCTGTNN
TGTCGCAGTATCTGTC




CCGGACACCCTATAT
||||||| |
ACACCCTATGTCGCA
|||||||
TGTCGCAGTATCTGTC
||| |||
TAT--GTCGCAGTATCTG

Image source: <http://ngm.nationalgeographic.com/your-shot/jigsaw-puzzles>

Analyzing reads



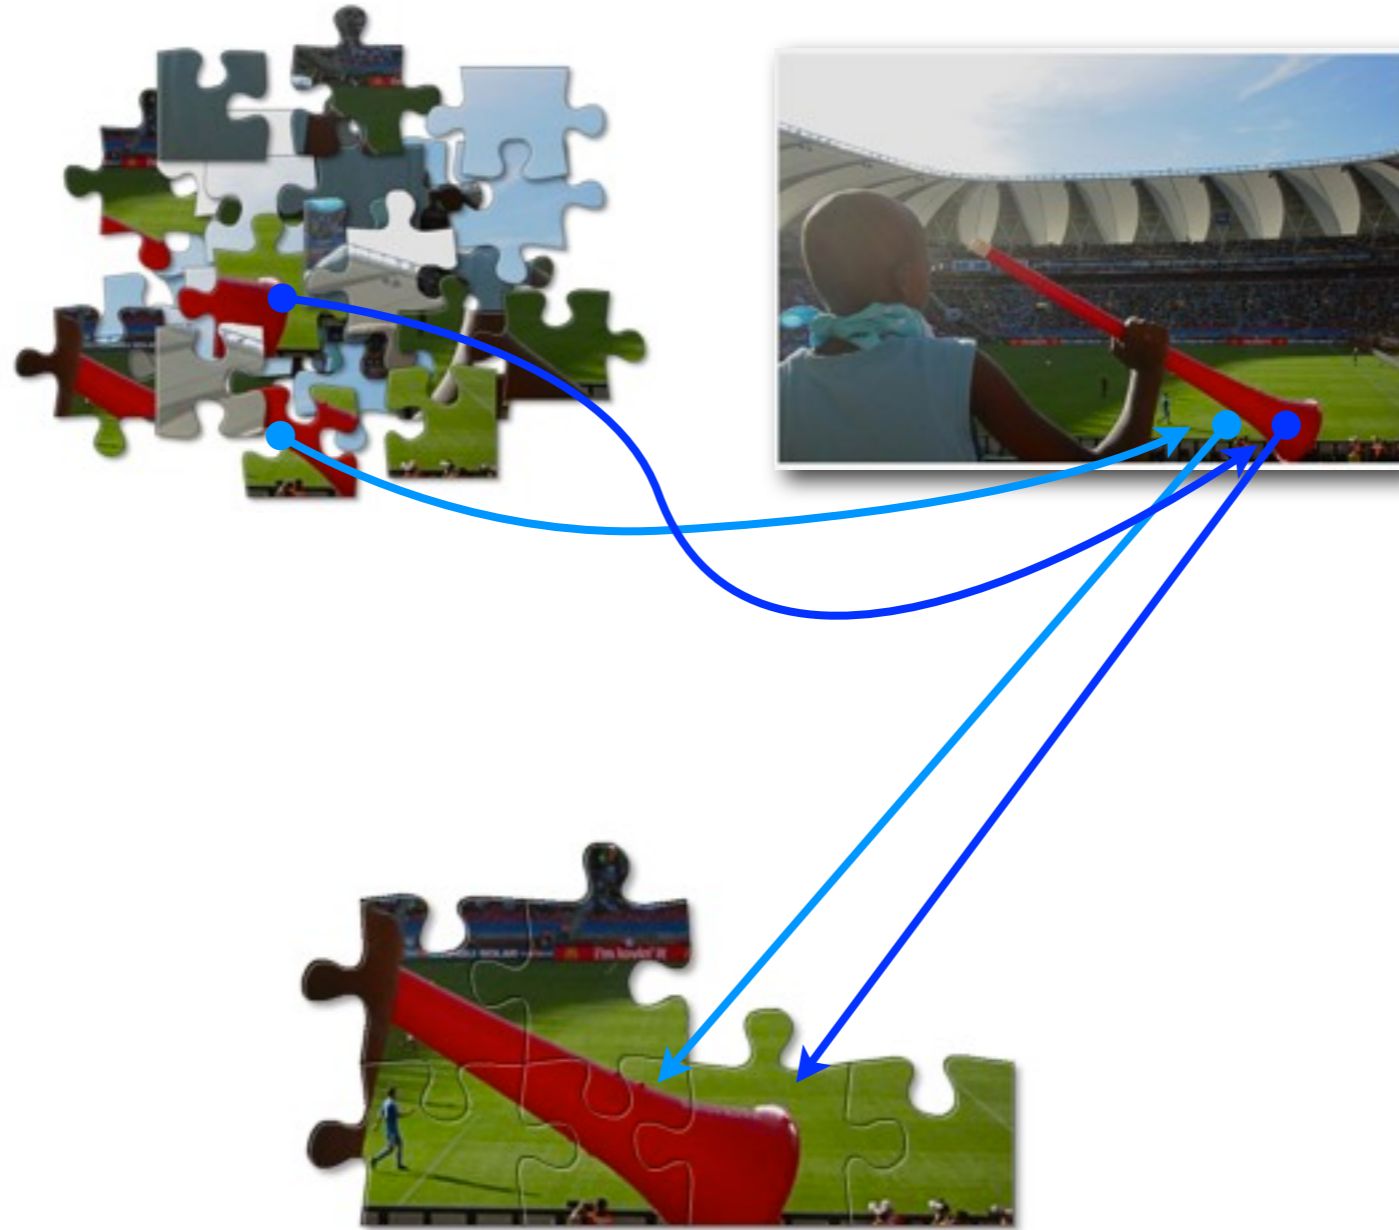
de novo



Comparative

Image source: <http://ngm.nationalgeographic.com/your-shot/jigsaw-puzzles>

Comparative



Comparative

Image source: <http://ngm.nationalgeographic.com/your-shot/jigsaw-puzzles>

Comparative

TATGTCGCAGTATCTT
TATGTCGCAGTATCTG
CCGGACACCCTATAT **GTCGCAGTATCTGTCT**
TATGTCGCAGTATCTT **GTCGCAGTATCTGTNN**
ACACCCTATGTCGCA
CCGGACACCCTATAT **TATGTCGCAGTATCTG**
CCGGACACCCTATAT **GTCGCAGTATCTGTNN**
TGTCGCAGTATCTGTC

```

>MT dna:chromosome chromosome:GRCh37:MT:1:16569:1
GATCACAGGTCTATCACCCCTATTAACCACTCACGGGAGCTCTCCATGCATTTGGTATTTT
CGTCTGGGGGTATGCACGCGATAGCATTGCGAGACGCTGGAGCCGGAGCACCCCTATGTC
GCAGTATCTGTCTTTGATTCTGCCTCATCCTATTATTTATCGCACCTACGTTCAATATT
ACAGGCGAACATACTTACTAAAGTGTGTTAATTAATTAATGCTTGTAGGACATAATAATA
ACAATTGAATGTCTGCACAGCCACTTTCCACACAGACATCATAACAAAAATTTCCACCA
AACCCCTCCCTCCCGCTTCTGGCCACAGCACTTAAACACATCTCTGCCAAACCCAAAA
ACAAAGAACCCTAACACCAGCCTAACCCAGATTTCAAATTTTATCTTTTGGCGGTATGCAC
TTTTAACAGTCAACCCCACTAACACATTATTTTCCCCTCCCACTCCATACTACTAAT
CTCATCAATACAACCCCGCCATCCTACCCAGCACACACACACCCGCTGCTAACCCATA
CCCCGAACCAACCAAACCCCAAAGACACCCCCACAGTTTATGTAGCTTACCTCCTCAA
GCAATACACTGACCCGCTCAAACCTCTGGATTTTGGATCCACCCAGCGCCTTGGCCTAAA
CTAGCCTTTCTATTAGCTCTTAGTAAGATTACACATGCAAGCATCCCCGTTCCAGTGAGT
TCACCTCTAAATCACCACGATCAAAAGGAACAAGCATCAAGCACGCAGCAATGCAGCTC
AAAACGCTTAGCCTAGCCACACCCCAACGCAAGCAGTGATTAACCTTTAGCAATAA
ACGAAAGTTAACTAAGCTATACTAAACCCAGGGTGGTCAATTTCTGTCAGCCACCGC
GGTCACACGATTAACCCAACTCAATAGAGCGTGGCGTAAAGAGTGTTTTAGATCACCCC
TCCCAATAAACCTAAACTCACCTGATTTGTAAAAACTCCAGTTGACACAAAATAGAC
TAGGAAAGTGGCTTAAACATATCTAACACAATAAGCTAAGACCCAAACTGGGATTAGA
TACCCCACTATGCTTAGCCCTAACCTCAAAGTTAAATCAACAAAACCTGCTCGCCAGAA
CACTACGAGCCACAGCTTAAACTCAAAGTACCTGGCGGTGCTTCATATCCCTCTAGAGG
AGCCTGTTCTTAATCGATAACCCCGAACAACCTCACCACTCTTGCTCAGCCTATATA
    
```

CCGGACACCCTATAT
 |||||
ACACCCTATGTCGCA
 |||||
TGTCGCAGTATCTGTC
 ||| |||
TAT--GTCGCAGTATCTG

Comparative

Comparative

TATGTCGCAGTATCTT
 TATGTCGCAGTATCTG
 CCGGACACCCTATAT GTCGCAGTATCTGTCT
 TATGTCGCAGTATCTT GTCGCAGTATCTGTNN
 ACACCCTATGTCGCA
 CCGGACACCCTATAT TATGTCGCAGTATCTG
 CCGGACACCCTATAT TGTGCGCAGTATCTGTC
 GTCGCAGTATCTGTNN
 TGTCGCAGTATCTGTC

>MT dna:chromosome:chromosome:GRCh37:MT:1:16569:1

TTT
 GTC
 ATT
 ATA
 CCA
 AAA
 AAC
 AAT
 ATA
 AAA
 AAA
 AGT
 CTC
 TAA
 CGC
 CCC
 GAC
 AGA
 TACCCCACTATGCTTAGCCCTAACCTCAACAGTTAAATCAACAAAAGTCTCGCCAGAA
 CACTACGAGCCACAAGTTAAAGTCMAAGGACTGCGCGTGCCTTAAAGTCTAGAGG
 AGCCTGTTCTGTAATCGATAACCCCGTCAACTCACCACTCTTCTCAGCCTATATA

**H. sapiens, Nature, 2000
and Science, 2000**

CCGGACACCCTATAT
 ||||| |
 ACACCCTATGTCGCA
 |||||
 TGTCGCAGTATCTGTC
 ||| ||
 TAT--GTCGCAGTATCTG

Comparative

Comparative


TATGTCGCAGTATCTT
 TATGTCGCAGTATCTG
 CCGGACACCCTATAT GTCGCAGTATCTGTCT
 TATGTCGCAGTATCTT GTCGCAGTATCTGTNN
 ACACCCTATGTCGCA
 CCGGACACCCTATAT TATGTCGCAGTATCTG
 CCGGACACCCTATAT GTCGCAGTATCTGTC
 GTCGCAGTATCTGTNN
 TGTCGCAGTATCTGTC



CCGGACACCCTATAT
 ||||| |
 ACACCCTATGTCGCA
 |||||
 TGTCGCAGTATCTGTC
 ||| |||
 TAT--GTCGCAGTATCTG

Comparative

Reference genome



Smith-Waterman

Aligning two sequences is a classic (and extremely important) problem in computational biology.

An 'efficient' solution is provided by the Smith-Waterman algorithm which produces the 'best' alignment under some statistical model.

It handles insertions and deletions elegantly (the default does not handle base qualities), but is too slow for short reads.

(`Biostrings::pairwiseAlignment()`)

Smith-Waterman

Aligning **d** reads of length **m** to reference of length **n** is $O(dmn)$

Say:

m = 100 nt

d = 2 billion (2×10^9) reads

n = 3 billion (3×10^9) nt \approx human

} \approx 1 week-long run of



Illumina HiSeq 2000

Source: http://www.illumina.com/systems/hiseq_2000.ilmn

Total of (6×10^{20}) Smith-Waterman cell updates required

A cluster of 1,000 6 Ghz processors, where each processor computes 1 cell update per clock cycle, would take >3 years

Alignment

Take a read:

CTCAAACCTCCTGACCTTTGGTGATCCACCCGCCTNGGCCTTC

And a reference sequence:

```
>MT dna:chromosome chromosome:GRCh37:MT:1:16569:1
GATCACAGGTCTATCACCTATTAACCACTCACGGGAGCTCTCCATGCATTTGGTATTTT
CGTCTGGGGGATGCACGCGATAGCATTGCGAGACGCTGGAGCCGGAGACCCTATGTC
GCAGTATCTGTCTTTGATTCCTGCCTCATCTATTATTTATCGCACCTACGTTCAATATT
ACAGGCGAACATACTTACTAAAGTGTGTTAATTAATTAATGCTTGTAGGACATAATAATA
ACAATTGAATGTCTGCACAGCCACTTTCCACACAGACATCATAACAAAAAATTTCCACCA
AACCCCTCCTCCCGCTTCTGGCCACAGCACTTAAACACATCTCTGCCAAACCCCAAAA
ACAAAGAACCCTAACACCAGCCTAACAGATTTCAAATTTTATCTTTTGGCGGTATGCAC
TTTTAACAGTCACCCCACTAACACATTTTCCCTCCCCTCCCATACTACTAAT
CTCATCAATACAACCCCGCCATCCTACCAGCACACACACACCGCTGCTAACCCATA
CCCCGAACCAACCAAACCCCAACAGACAGCCGCGACACTTTATCTAGCTTACCTCCTCAA
GCAATACACTGACCGCTCAAACCTCCTGGATTTTGGATCCACCCAGCGCCTTGGCCTAA
CTAGCCTTTCTATTAGCTTTAGTAAAGATTACATGCAAGCATCCCTCCCTCCAGTCAAT
TCACCCTCTAATCACCACGATCAAAAGGAACAAGCATCAAGCACGCAGCAATGCAGCTC
AAAACGCTTAGCCTAGCCACACCCCCACGGGAAACAGCAGTGATTAACCTTTAGCAATAA
ACGAAAGTTTAACTAAGCTATACTAACCAGGGTGGTCAATTTTCGTGCCAGCCACCGC
GGTACACGATTAACCAAGTCAATAGAAGCCGGCGTAAAGAGTGTTTTAGATCACCCCC
TCCCAATAAAGCTAAAACCTCACCTGAGTTGTAATAAAGTCCAGTTGACACAAAATAGAC
TACGAAAGTGGCTTTAACATATCTGAACACACAATAGCTAAGACCCAAACTGGGATTAGA
TACCCACTATGCTTAGCCCTAACCTCAACAGTTAAATCAACAAAAGTCTCGCCAGAA
CACTACGAGCCACAGCTTAAAACCTCAAAGGACCTGGCGGTGCTTCATATCCCTCTAGAGG
AGCCTGTTCTGTAATCGATAAACCCCGATCAACCTCACCACTCTTGCTCAGCCTATATA
CCGCCATCTTCAGCAAACCCTGATGAAGGCTACAAAGTAAGCGCAAGTACCCACGTAAAG
ACGTTAGGTCAAGGTGTAGCCCATGAGGTGGCAAGAAATGGGCTACATTTTCTACCCAG
AAAACCTACGATAGCCCTTATGAACTTAAGGGTCAAGGTGGATTTAGCAGTAACTAAG
AGTAGAGTGCTTAGTTGAACAGGGCCCTGAAGCGGTACACACCAGCCCGTCAACCTCCTC
AAGTATACTCAAAGGAGATTAACTAAACCCCTAGCCATTATATAGAGGAGACAAGT
CGTAACTCAAACCTCCTGCCTTTGGTGATCCACCCGCCTTGGCCTACCTGCATAATGAAG
AAGCACCAACTTACACTTAGGAGATTTCACCTAACTTACCCCTCTGAGCTAAACCTA
GCCCAAACCCACTCCACCTTACTACCAGACAACCTTAGCCAAACCATTTACCCAAATAA
AGTATAGGCGATAGAAATTGAAACCTGGCGCAATAGATATAGTACCGCAAGGAAAGATG
AAAAATTATAACCAAGCATAATATAGCAAGGACTAACCCCTATACCTTCTGCATAATGAA
TTAACTAGAAATAACTTTGCAAGGAGAGCCAAAGCTAAGACCCCGAAACAGACGAGCT
ACCTAAGAAGAGCTAAAAGAGCACACCCGCTATGTAGCAAATAGTGGGAAGATTTATA
GGTAGAGGCGACAAACCTACCGAGCCTGGTGATAGCTGGTTGTCCAAGATAGAATCTTAG
TTCAACTTTAAATTTGCCACAGAACCTCTAAATCCCCTTGTAATTTAACTGTTAGTC
CAAAGAGGAACAGCTCTTTGGACACTAGGAAAAAACCCTTGTAAGAGAGGTAATAAATTA
ACACCCATAGTAGGCCTAAAAGCAGCCACCAATTAAGAAAGCGTTCAAGCTCAACACCCA
CTACCTAAAAAATCCCAAACATATACTGAACCTCACACCCAATTGGACCAATCTATC
ACCTATAGAAGAATAACTAATGTTAGTATAAGTAAACATGAAAACATTCTCCTCCGCATAAGC
```

How do we determine the read's point of origin with respect to the reference?

Match 1:



Match 2:



Which match is better?

Say match 2 is correct. Why are there still mismatches and gaps?

Alignment

Take a read:

CTCAAACCTCCTGACCTTTGGTGATCCA

And a reference sequence:

```
>MT dna:chromosome chromosome:GRCh37:MT:1:16569:1
GATCACAGGTCTATCACCTATTAACCACTCACGGGAGCTCTCCATGCATTTGGTATTTT
CGTCTGGGGGGTATGCACGCGATAGCATTGCGAGACGCTGGAGCCGGAGCACCCCTATGTC
GCAGTATCTGTCTTTGATTCTGCCTCATCCTATTATTTATCGCACCTACGTTCAATATT
ACAGGCGAACATACTTACTAAAGTGTGTTAATTAATTAATGCTTGTAGGACATAATAATA
ACAATTGAATGTCTGCACAGCCACTTTCCACACAGACATCATAACAAAAATTTCCACCA
AACCCCTCCCGCTTCTGGCCACAGCACTTAAACACATCTCTGCCAAACCCCAAAA
ACAAAGAACCCTAACACCAGCCTAACAGATTTCAAATTTTATCTTTTGGCGGTATGCAC
TTTTAACAGTCACCCCACTAACACATTATTTTCCCCTCCCACTCCCACTACTACTAAT
CTCATCAATAACAACCCCGCCATCCTACCCAGCACACACACACCGCTGCTAACCCATA
CCCCGAACCAACCAACCCCAAGACACCCCCACAGTTTATGTAGCTTACCTCCTCAA
GCAATACACTGACCCGCTCAAACCTCTGGATTTTGTGATCCACCAGCGCCTTGGCCTAA
CTAGCCTTTCTATTAGCTCTTAGTAAGATTACACATGCAAGCATCCCCGTTCCAGTGAGT
TCACCCTCTAAATCACCACGATCAAAAGGAACAAGCATCAAGCACGCAGCAATGCAGCTC
AAAACGCTTAGCCTAGCCACACCCCCACGGGAAACAGCAGTGATTAACCTTTAGCAATAA
ACGAAAGTTTAACTAAGCTATACTAACCCAGGGTTGGTCAATTTCTGTGCCAGCCACCGC
GGTCACACGATTAACCCAAGTCAATAGAAGCCGGCGTAAAGAGTGTTTTAGATCACCC
TCCCAATAAAGCTAAAACCTCACCTGAGTTGTAAAAACTCCAGTTGACACAAAATAGAC
TACGAAAGTGGCTTTAACATATCTGAACACACAATAGCTAAGACCCAAACTGGGATTAGA
TACCCCACTATGCTTAGCCCTAACCTCAACAGTTAAATCAACAAAACCTGCTCGCCAGAA
CACTACGAGCCACAGCTTAAAACCTCAAAGGACCTGGCGGTGCTTCATATCCCTCTAGAGG
AGCCTGTTCTGTAATCGATAAACCCGATCAACCTCACCACTCTTGCTCAGCCTATATA
CCGCCATCTTCAGCAAACCTGATGAAGGCTACAAAGTAAGCGCAAGTACCCACGTAAG
ACGTTAGGTCAAGGTGTAGCCCATGAGGTGGCAAGAAATGGGCTACATTTTCTACCCAG
AAAACCTACGATAGCCCTTATGAACTTAAGGGTCAAGGTGGATTTAGCAGTAACTAAG
AGTAGAGTGCTTAGTTGAACAGGGCCCTGAAGCGCGTACACACCGCCCGTACCCTCCTC
AAGTATACTTCAAAGGACATTTAACTAAAACCCCTACGCATTTATATAGAGGAGACAAGT
CGTAACTCAAACCTCTGGCCTTTGGTGATCCACCCGCTTGGCCTACCTGCATAATGAA
AAGCACCCAACTTACACTTAGGAGATTTCAACTTAACTTGACCGCTCTGAGCTAAACCTA
GCCCCAAACCCACTCCACCTTACTACCAGACAACCTTAGCCAAACCTTTACCCAAATAA
AGTATAGGCGATAGAAATTGAAACCTGGCGCAATAGATATAGTACCGCAAGGGAAAGATG
AAAAATTATAACCAAGCATAATATAGCAAGGACTAACCCCTATACCTTCTGCATAATGAA
TTAACTAGAAATAACTTTGCAAGGAGAGCCAAAGCTAAGACCCCGAAACAGACGAGCT
ACCTAAGAACAGCTAAAAGAGCACACCCGCTATGTAGCAAATAGTGGGAAGATTTATA
GGTAGAGGCGACAAACCTACCGAGCCTGGTGATAGCTGGTTGTCCAAGATAGAATCTTAG
TTCAACTTTAAATTTGCCACAGAACCTCTAAATCCCCTTGTAATTTAACTGTTAGTC
CAAAGAGGAACAGCTCTTTGGACACTAGGAAAAACCTTGTAAGAGAGAGTAAAAATTTA
ACACCCATAGTAGGCCTAAAAGCAGCCACCAATTAAGAAAGCGTTCAAGCTCAACACCCA
CTACCTAAAAAATCCCAAACATATAACTGAACTCCTCACACCCAATTGGACCAATCTATC
ACCCTATAGAAGAACTAATGTTAGTATAAGTAACATGAAAACATTCTCCTCCGCATAAGC
```

Which match is better?

Match 1:

```
Read
CTCAAACCTCCTGACCTTTGGTGATCCA
|||||
Reference
CTCAAACCTCCTGACCTTTGGTGATCCA
```

Match 2:

```
Read
CTCAAACCTCCTGACCTTTGGTGATCCA
|||||
Reference
CTCAAACCTCCTGACCTTTGGTGATCCA
```

Is there any way to break the tie?

Two types of qualities

- **Base (sequence) quality**

Represents the chance that the sequence machine made an error. Produced by the sequence machine (possibly with some post-processing, “calibration”). The ‘Q’ in FASTQ files.

- **Alignment quality**

Represents the chance that the alignment is wrong. Produced by the alignment software.

Does base quality really reflect the chance of a sequence error?

Alignment

Take a read:

CTCAAACCTCCTGACCTTTGGTGATCCA

And a reference sequence:

```
>MT dna:chromosome chromosome:GRCh37:MT:1:16569:1
GATCACAGGTCTATCACCTATTAACCACTCACGGGAGCTCTCCATGCATTTGGTATTTT
CGTCTGGGGGGTATGCACGCGATAGCATTGCGAGACGCTGGAGCCGGAGCACCCCTATGTC
GCAGTATCTGTCTTTGATTCTGCCTCATCCTATTATTTATCGCACCTACGTTCAATATT
ACAGGCGAACATACTTACTAAAGTGTGTTAATTAATTAATGCTTGTAGGACATAATAATA
ACAATTGAATGTCTGCACAGCCACTTTCCACACAGACATCATAACAAAAATTTCCACCA
AACCCCTCCCGCTTCTGGCCACAGCACTTAAACACATCTCTGCCAAACCCCAAAA
ACAAAGAACCCTAACACCAGCCTAACAGATTTCAAATTTTATCTTTTGGCGGTATGCAC
TTTTAACAGTCACCCCACTAACACATTATTTCCCTCCCACTCCATACTACTAAT
CTCATCAATAACAACCCCGCCATCCTACCCAGCACACACACACCGCTGCTAACCCATA
CCCCGAACCAACCAACCCCAAGACACCCCCACAGTTTATGTAGCTTACCTCCTCAA
GCAATACACTGACCCGCTCAAACCTCTGGATTTTGTGATCCACCAGCGCCTTGGCCTAA
CTAGCCTTTCTATTAGCTCTTAGTAAGATTACACATGCAAGCATCCCCGTTCCAGTGAGT
TCACCCTCTAAATCACCACGATCAAAAGGAACAAGCATCAAGCACGCAGCAATGCAGCTC
AAAACGCTTAGCCTAGCCACACCCCCACGGGAAACAGCAGTGATTAACCTTTAGCAATAA
ACGAAAGTTTAACTAAGCTATACTAACCCAGGGTTGGTCAATTTCTGTGCCAGCCACCGC
GGTCACACGATTAACCCAAGTCAATAGAAGCCGGCGTAAAGAGTGTTTTAGATCACCC
TCCCAATAAAGCTAAAACCTCACCTGAGTTGTAAAAAATCCAGTTGACACAAAATAGAC
TACGAAAGTGGCTTTAACATATCTGAACACACAATAGCTAAGACCCAAACTGGGATTAGA
TACCCCACTATGCTTAGCCCTAACCTCAACAGTTAAATCAACAAAATGCTCGCCAGAA
CACTACGAGCCACAGCTTAAAACCTCAAAGGACCTGGCGGTGCTTCATATCCCTCTAGAGG
AGCCTGTTCTGTAATCGATAAACCCGATCAACCTCACCACTCTTGCTCAGCCTATATA
CCGCCATCTTCAGCAAACCTGATGAAGGCTACAAAGTAAGCGCAAGTACCCACGTAAG
ACGTTAGGTCAAGGTGTAGCCCATGAGGTGGCAAGAAATGGGCTACATTTTCTACCCAG
AAAACCTACGATAGCCCTTATGAACTTAAGGGTCAAGGTGGATTTAGCAGTAACTAAG
AGTAGAGTGCTTAGTTGAACAGGGCCCTGAAGCGCGTACACACCGCCCGTACCCTCCTC
AAGTATACTTCAAAGGACATTTAACTAAAACCCCTACGCATTTATATAGAGGAGACAAGT
CGTAACTCAAACCTCTGGCCTTTGGTGATCCACCCGCTTGGCCTACCTGCATAATGAA
AAGCACCCAACTTACACTTAGGAGATTTCAACTTAACTTGACCGCTCTGAGCTAAACCTA
GCCCAAAACCCACTCCACCTTACTACCAGACAACCTTAGCCAAACCATTTACCCAATAA
AGTATAGGCGATAGAAATTGAAACCTGGCGCAATAGATATAGTACCGCAAGGGAAAGATG
AAAAATTATAACCAAGCATAATATAGCAAGGACTAACCCCTATACCTTCTGCATAATGAA
TTAACTAGAAATAACTTTGCAAGGAGAGCCAAAGCTAAGACCCCGAAACAGACGAGCT
ACCTAAGAACAGCTAAAAGAGCACACCCGCTATGTAGCAAATAGTGGGAAGATTTATA
GGTAGAGGCGACAAACCTACCGAGCCTGGTGATAGCTGGTTGTCCAAGATAGAATCTTAG
TTCAACTTTAAATTTGCCACAGAACCCTCTAAATCCCCTTGTAATTTAACTGTTAGTC
CAAAGAGGAACAGCTCTTTGGACACTAGGAAAAAACCTTGTAGAGAGAGTAAAAATTTA
ACACCCATAGTAGGCCTAAAAGCAGCCACCAATTAAGAAAGCGTTCAAGCTCAACACCCA
CTACCTAAAAAATCCCAAACATATAACTGAACTCCTCACACCCAATTGGACCAATCTATC
ACCCTATAGAAGAACTAATGTTAGTATAAGTAACATGAAAACATTCTCCTCCGCATAAGC
```

Which match is better?

Match 1:

Q=30

Read

```
CTCAAACCTCCTGACCTTTGGTGATCCA
|||||
CTCAAACCTCCTGACCTTTGGTGATCCA
```

Reference

Match 2:

Q=10

Read

```
CTCAAACCTCCTGACCTTTGGTGATCCA
|||||
CTCAAACCTCCTGACCTTTGGTGATCCA
```

Reference



Alignment

Read 1:

Best match:

Read

```
AGCTTATATGCTTTTCAGAGCGATACTAAAACCNAACTCA
|||||
AGCTTATATGCTTTTCAGAGCGATACTAAAACCTAACCTCA
Reference
```

Second-best match:

Read

```
AGCTTATATGCTATTTTCAGAGCGATACTAAAACCNAACTTA
|||||
AGCTTATATGCT-TTTCAGAGCGATACTAAAACCTAACCTCA
Reference
```

Read 2:

Best match:

Read

```
CTCAAACCTCCTGACCTTTGGTGATCCACCCGCCTNGGCCTTC
|||||
CTCAAACCTCCTG---TTTGGTGATCCACCCGCCTTGGCCTAC
Reference
```

Second-best match:

Read

```
CTCAAAGACCTGACCTTTGGTGATAAACCC-----GCCTNGGCCTTC
|||||
CTCA-----CCTGATTTTG--GATCCGCCAGCTGGCCTTGGCCTAA
Reference
```

For which read are we more confident that the best match is correct?

Alignment and Bioconductor

These days, most alignment is done using external tools.

However, it is worth knowing about

`matchPDict`

`matchPWM`

`pairwiseAlignment`

in Biostrings.

Popular aligners

- Bowtie
- BWA(-SW)
- MAQ
- SOAP2
- Novoalign
-

Many programs support more than one alignment 'mode' depending on command line settings.

The choice of settings is often unclear.

Which aligner is best?

- Two issues: (1) which aligner is the best implementation of a given policy? and (2) which policy is best?
- There has been surprisingly little investigation of which policy is best on real data. It is a hard problem.
- Most aligners have been evaluated in terms of **speed** and **completeness** (% of reads mapped).
- Completeness is probably the wrong metric.
- Some evaluation on simulated data, but we need more.
- Different aligners (policies) produce different end results, sometimes dramatically different.
- Answer also depends on “for what”.

Fileformats

- Input

FASTQ, FASTA, QSEQ, SFF

Vendor specific formats (like CSFASTA+QUAL)

- Output

BAM/SAM, program-specific format

Tip: Learn the UNIX shell, especially piping

```
gunzip -c INPUT.fastq.gz | \  
bowtie -m 1 -v 2 -p 4 -y --trim3 10 hsapiens_hg19 - | \  
gzip -c > OUTPUT.bwt.gz
```



```
@HWI-EAS146:5:1:1:961#0/1  
TCCGAGGCCAACCGAGGCTCCGCGGCGCTGNNNNNNNNNNCNNNNN  
+  
BBBB>A?B@;@BBBBBAA=BA=A%%%%%%%%%  
@HWI-EAS146:5:1:1:1595#0/1  
TCAGGAAGCAGGAAGAGCTGGTGCAGCAGGNNNNNNNNNNGNNNNN
```

name
sequence
quality scores

Interlude

Now for some perspectives on aligning RNA-seq data.

Junction reads

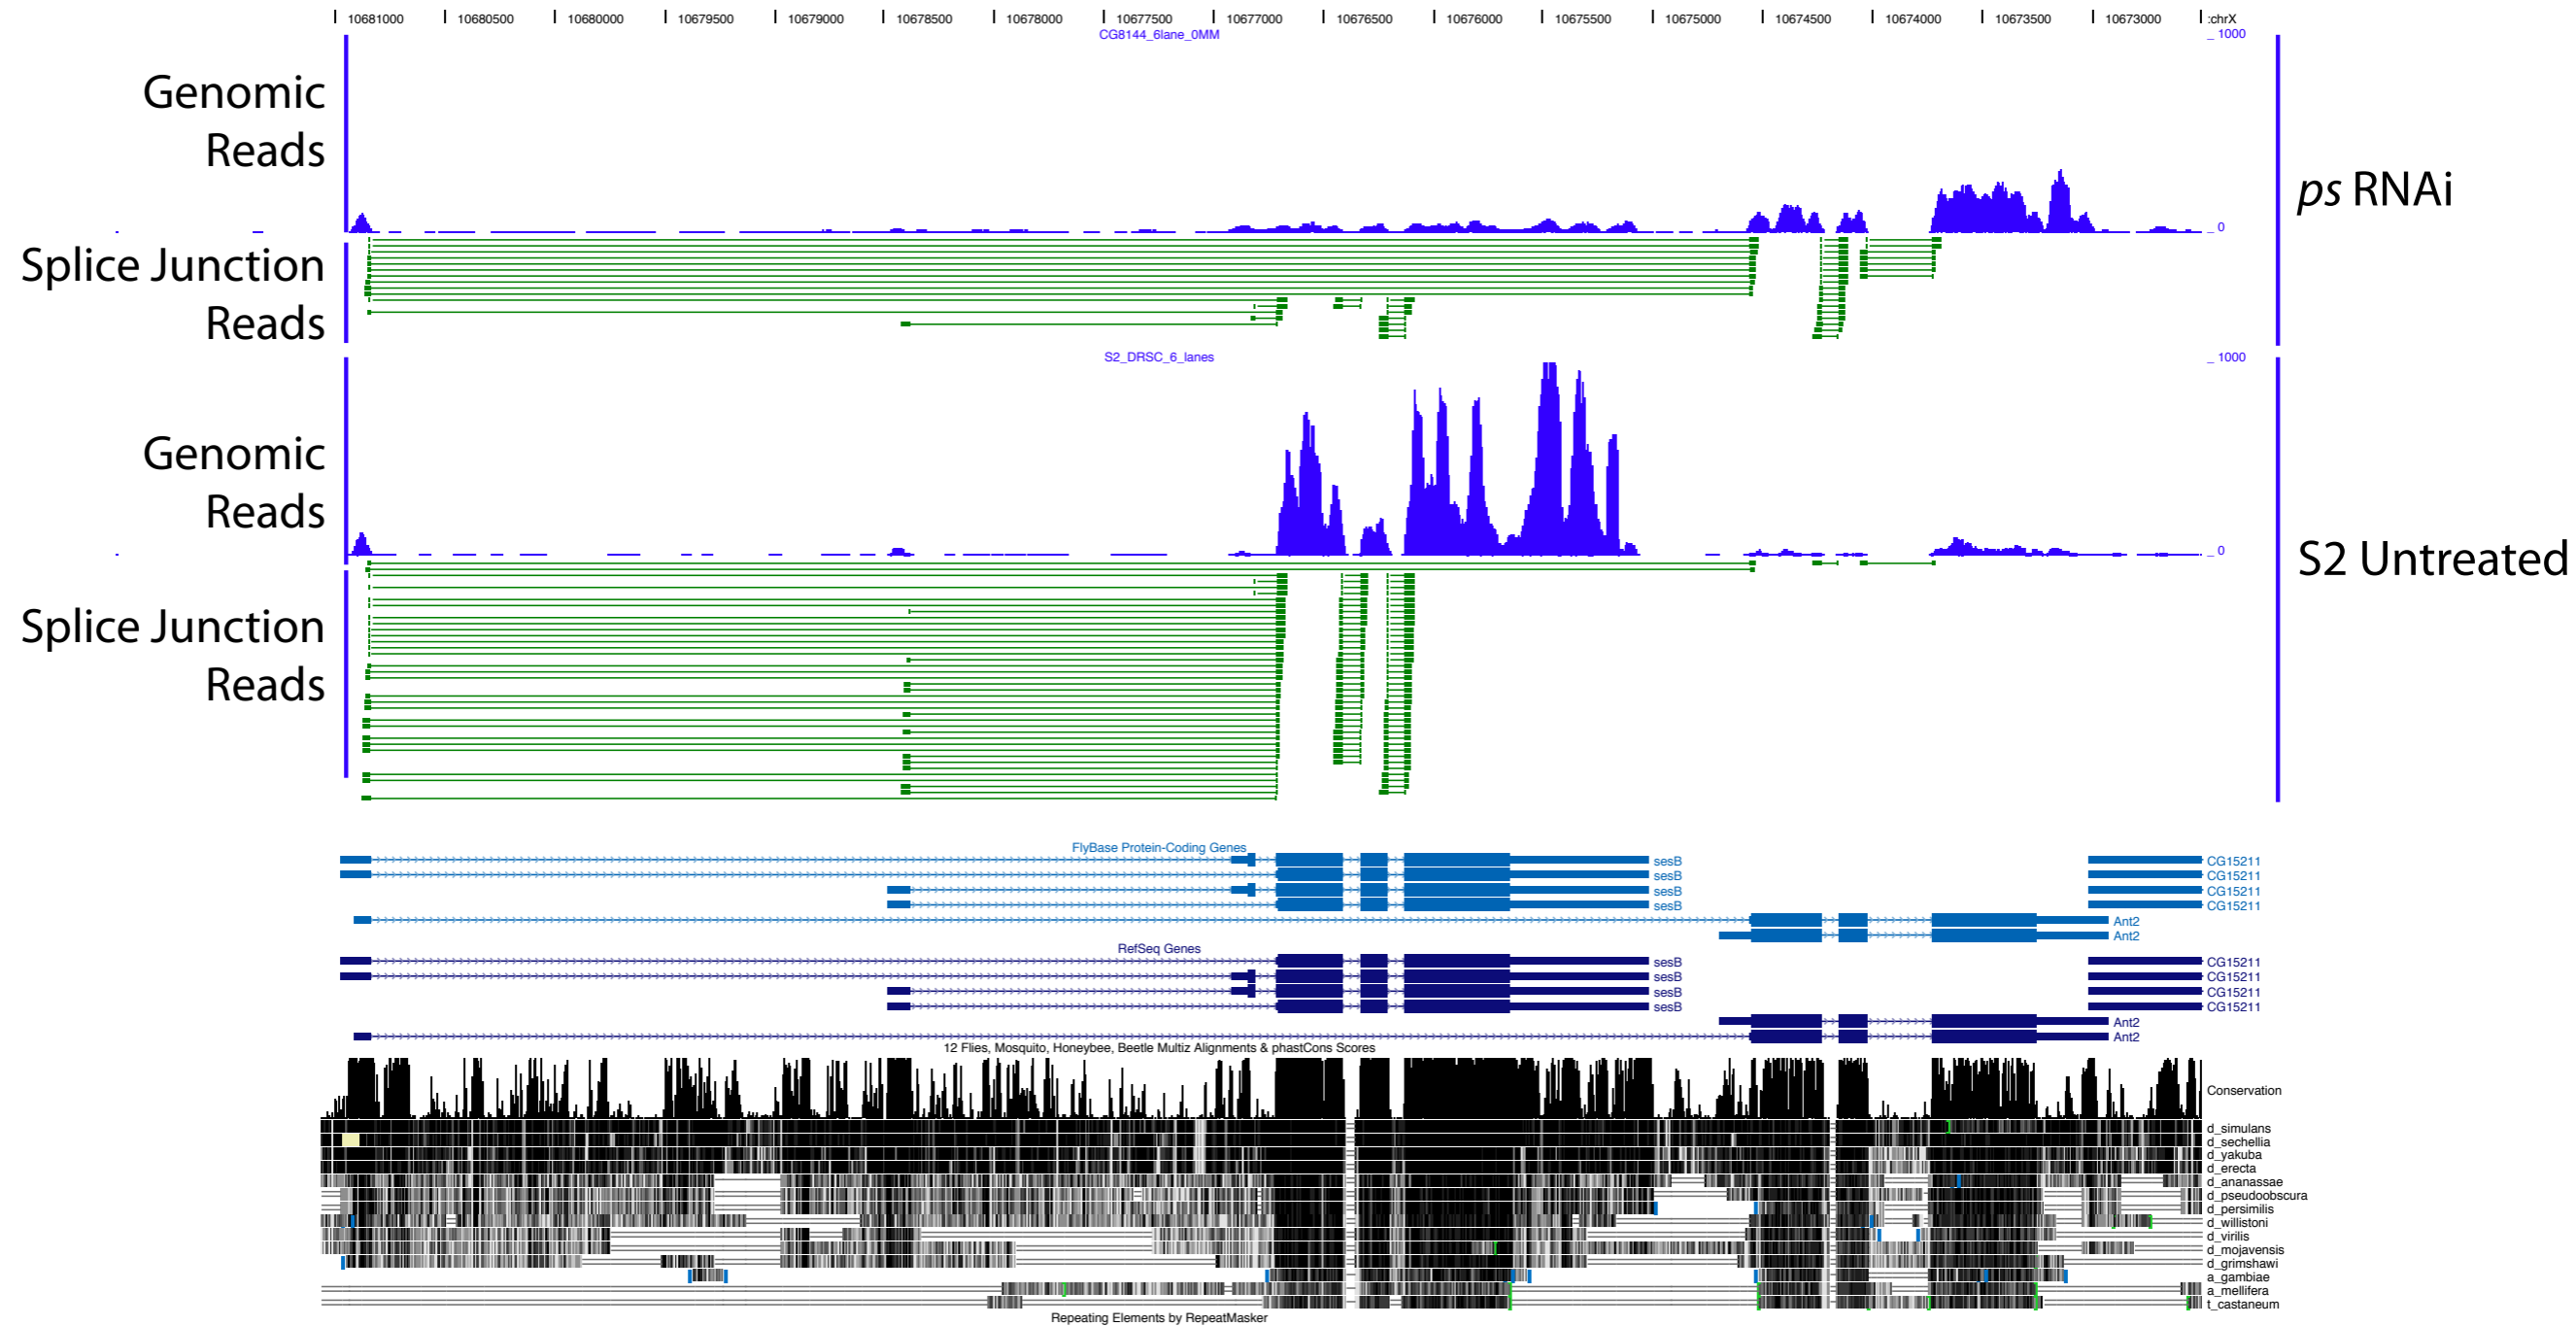


Image from Brenton Gravely



JOHNS HOPKINS
BLOOMBERG
SCHOOL OF PUBLIC HEALTH

Junction reads, zoom

S2_DRSC_6_lanes



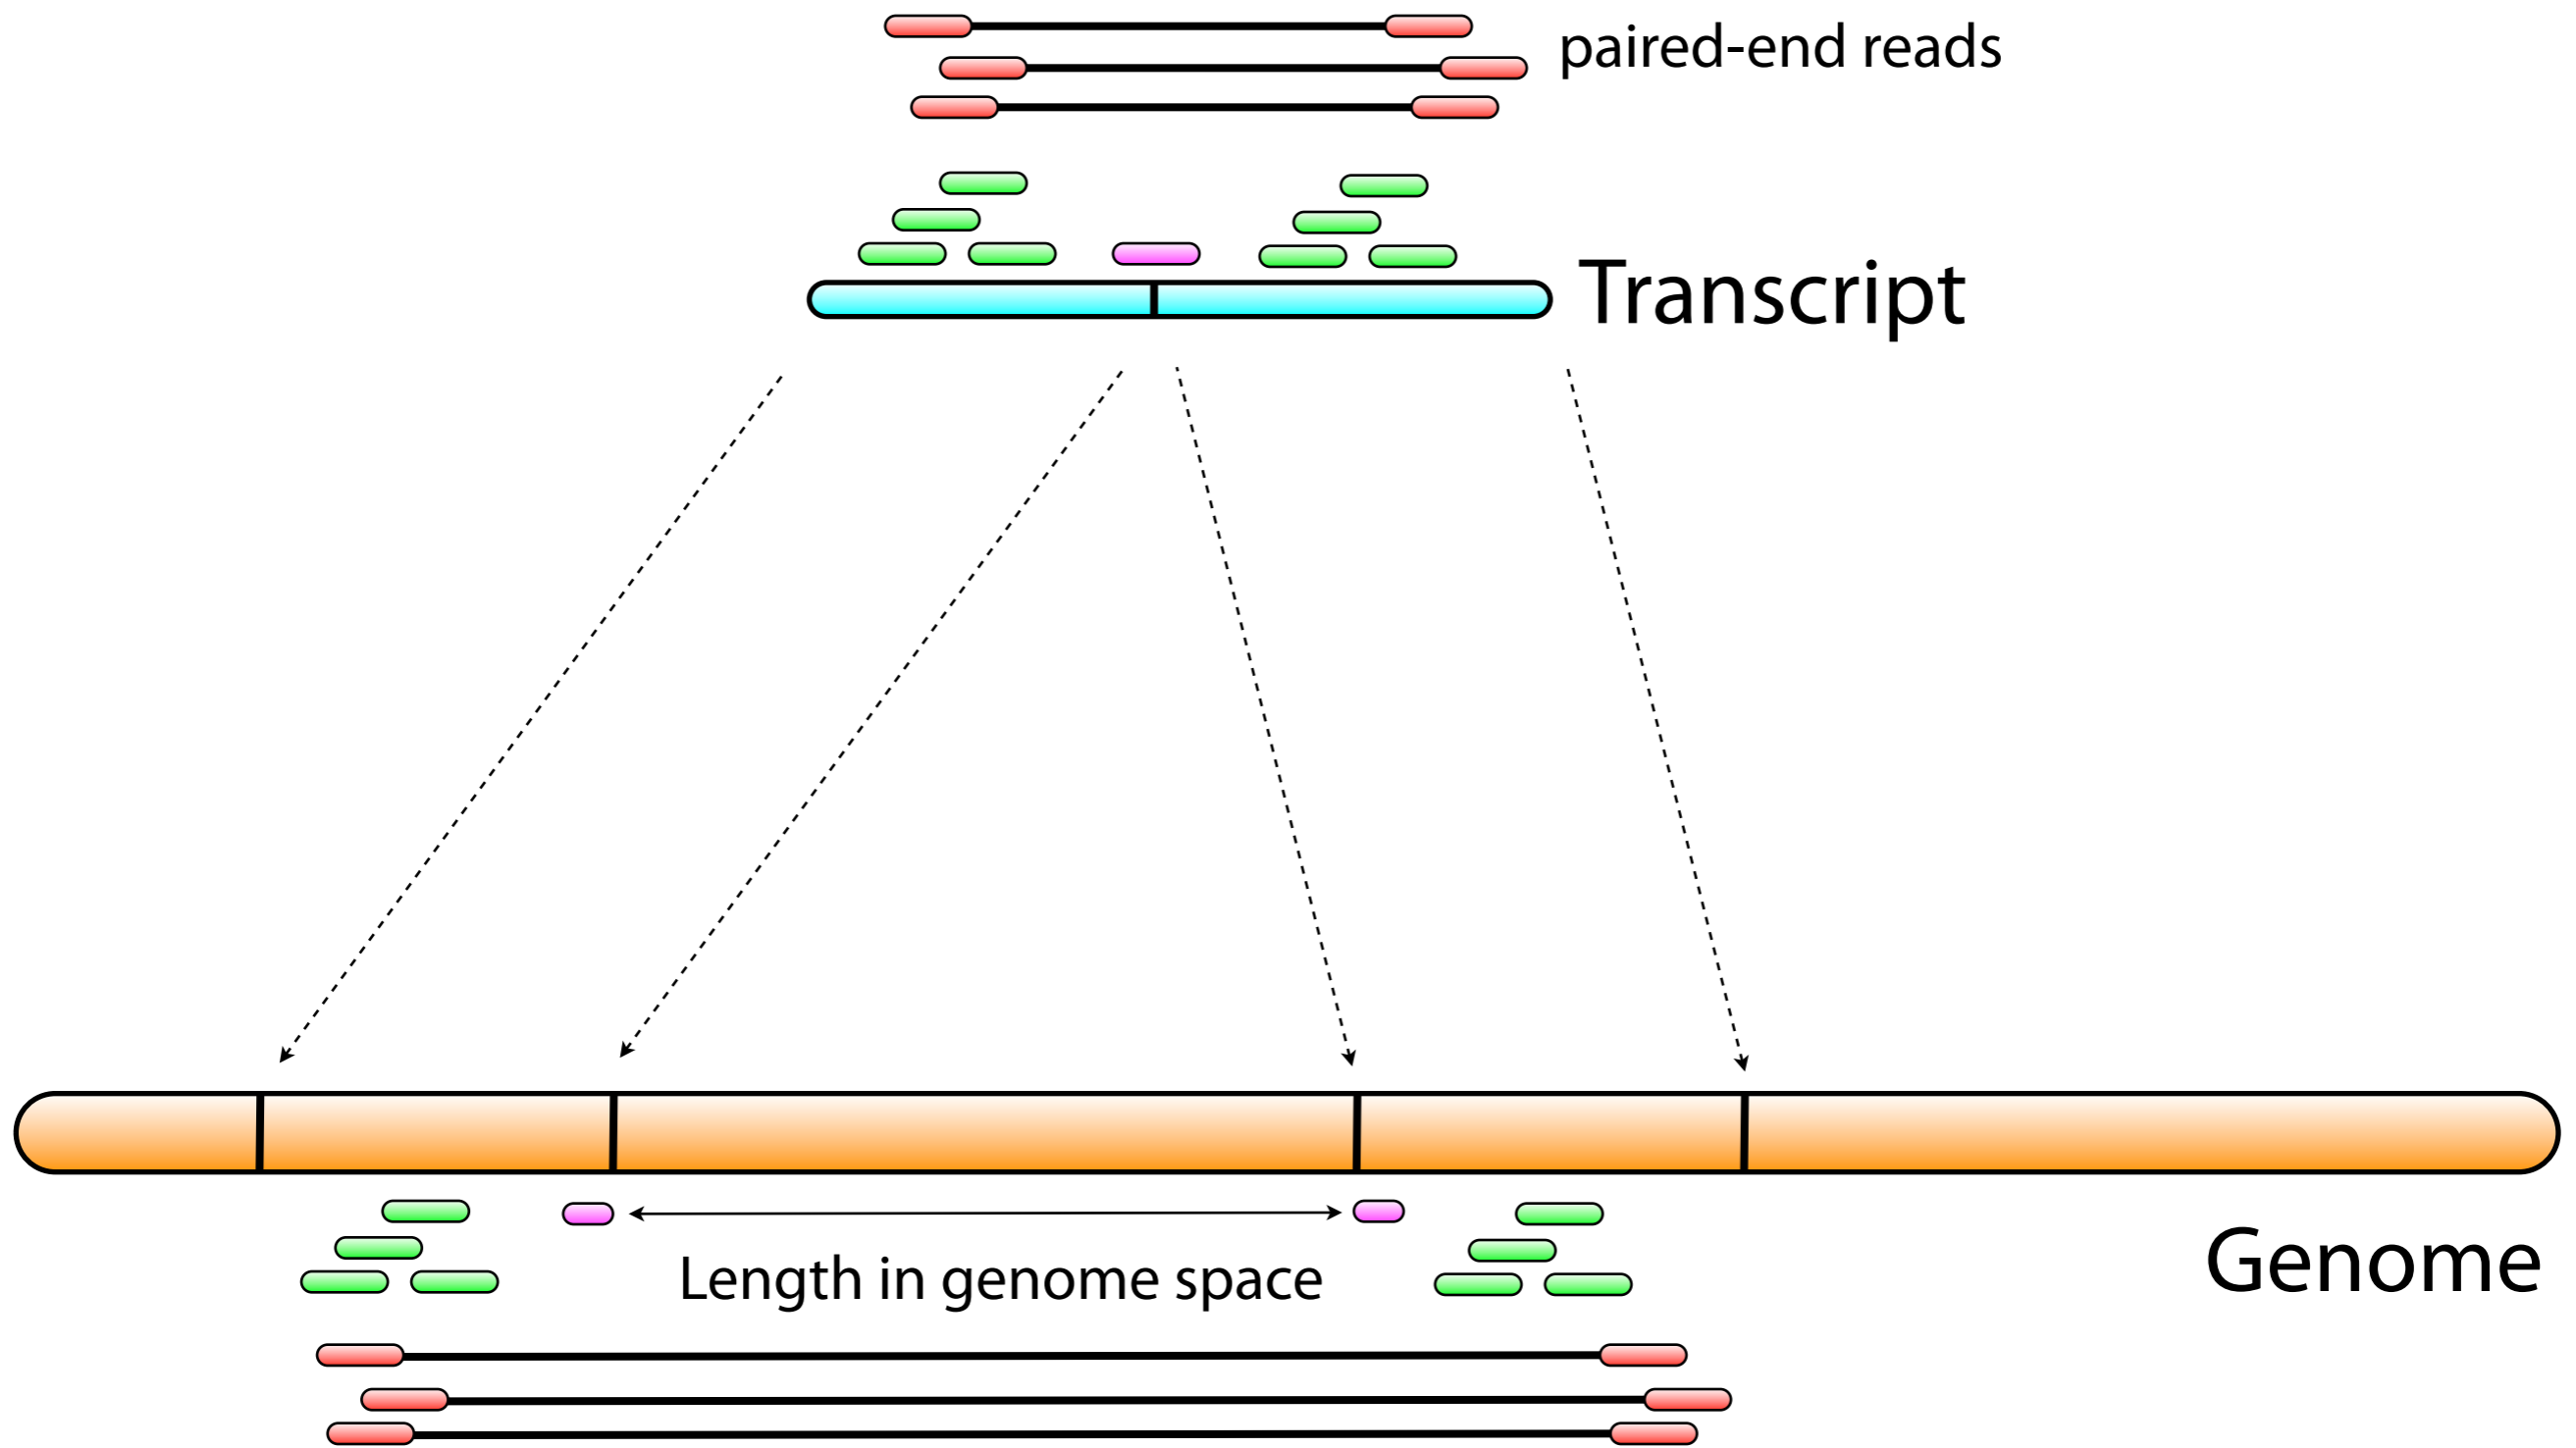


Image from Brenton Gravely

Mapping transcripts



Mapping reads to the transcriptome

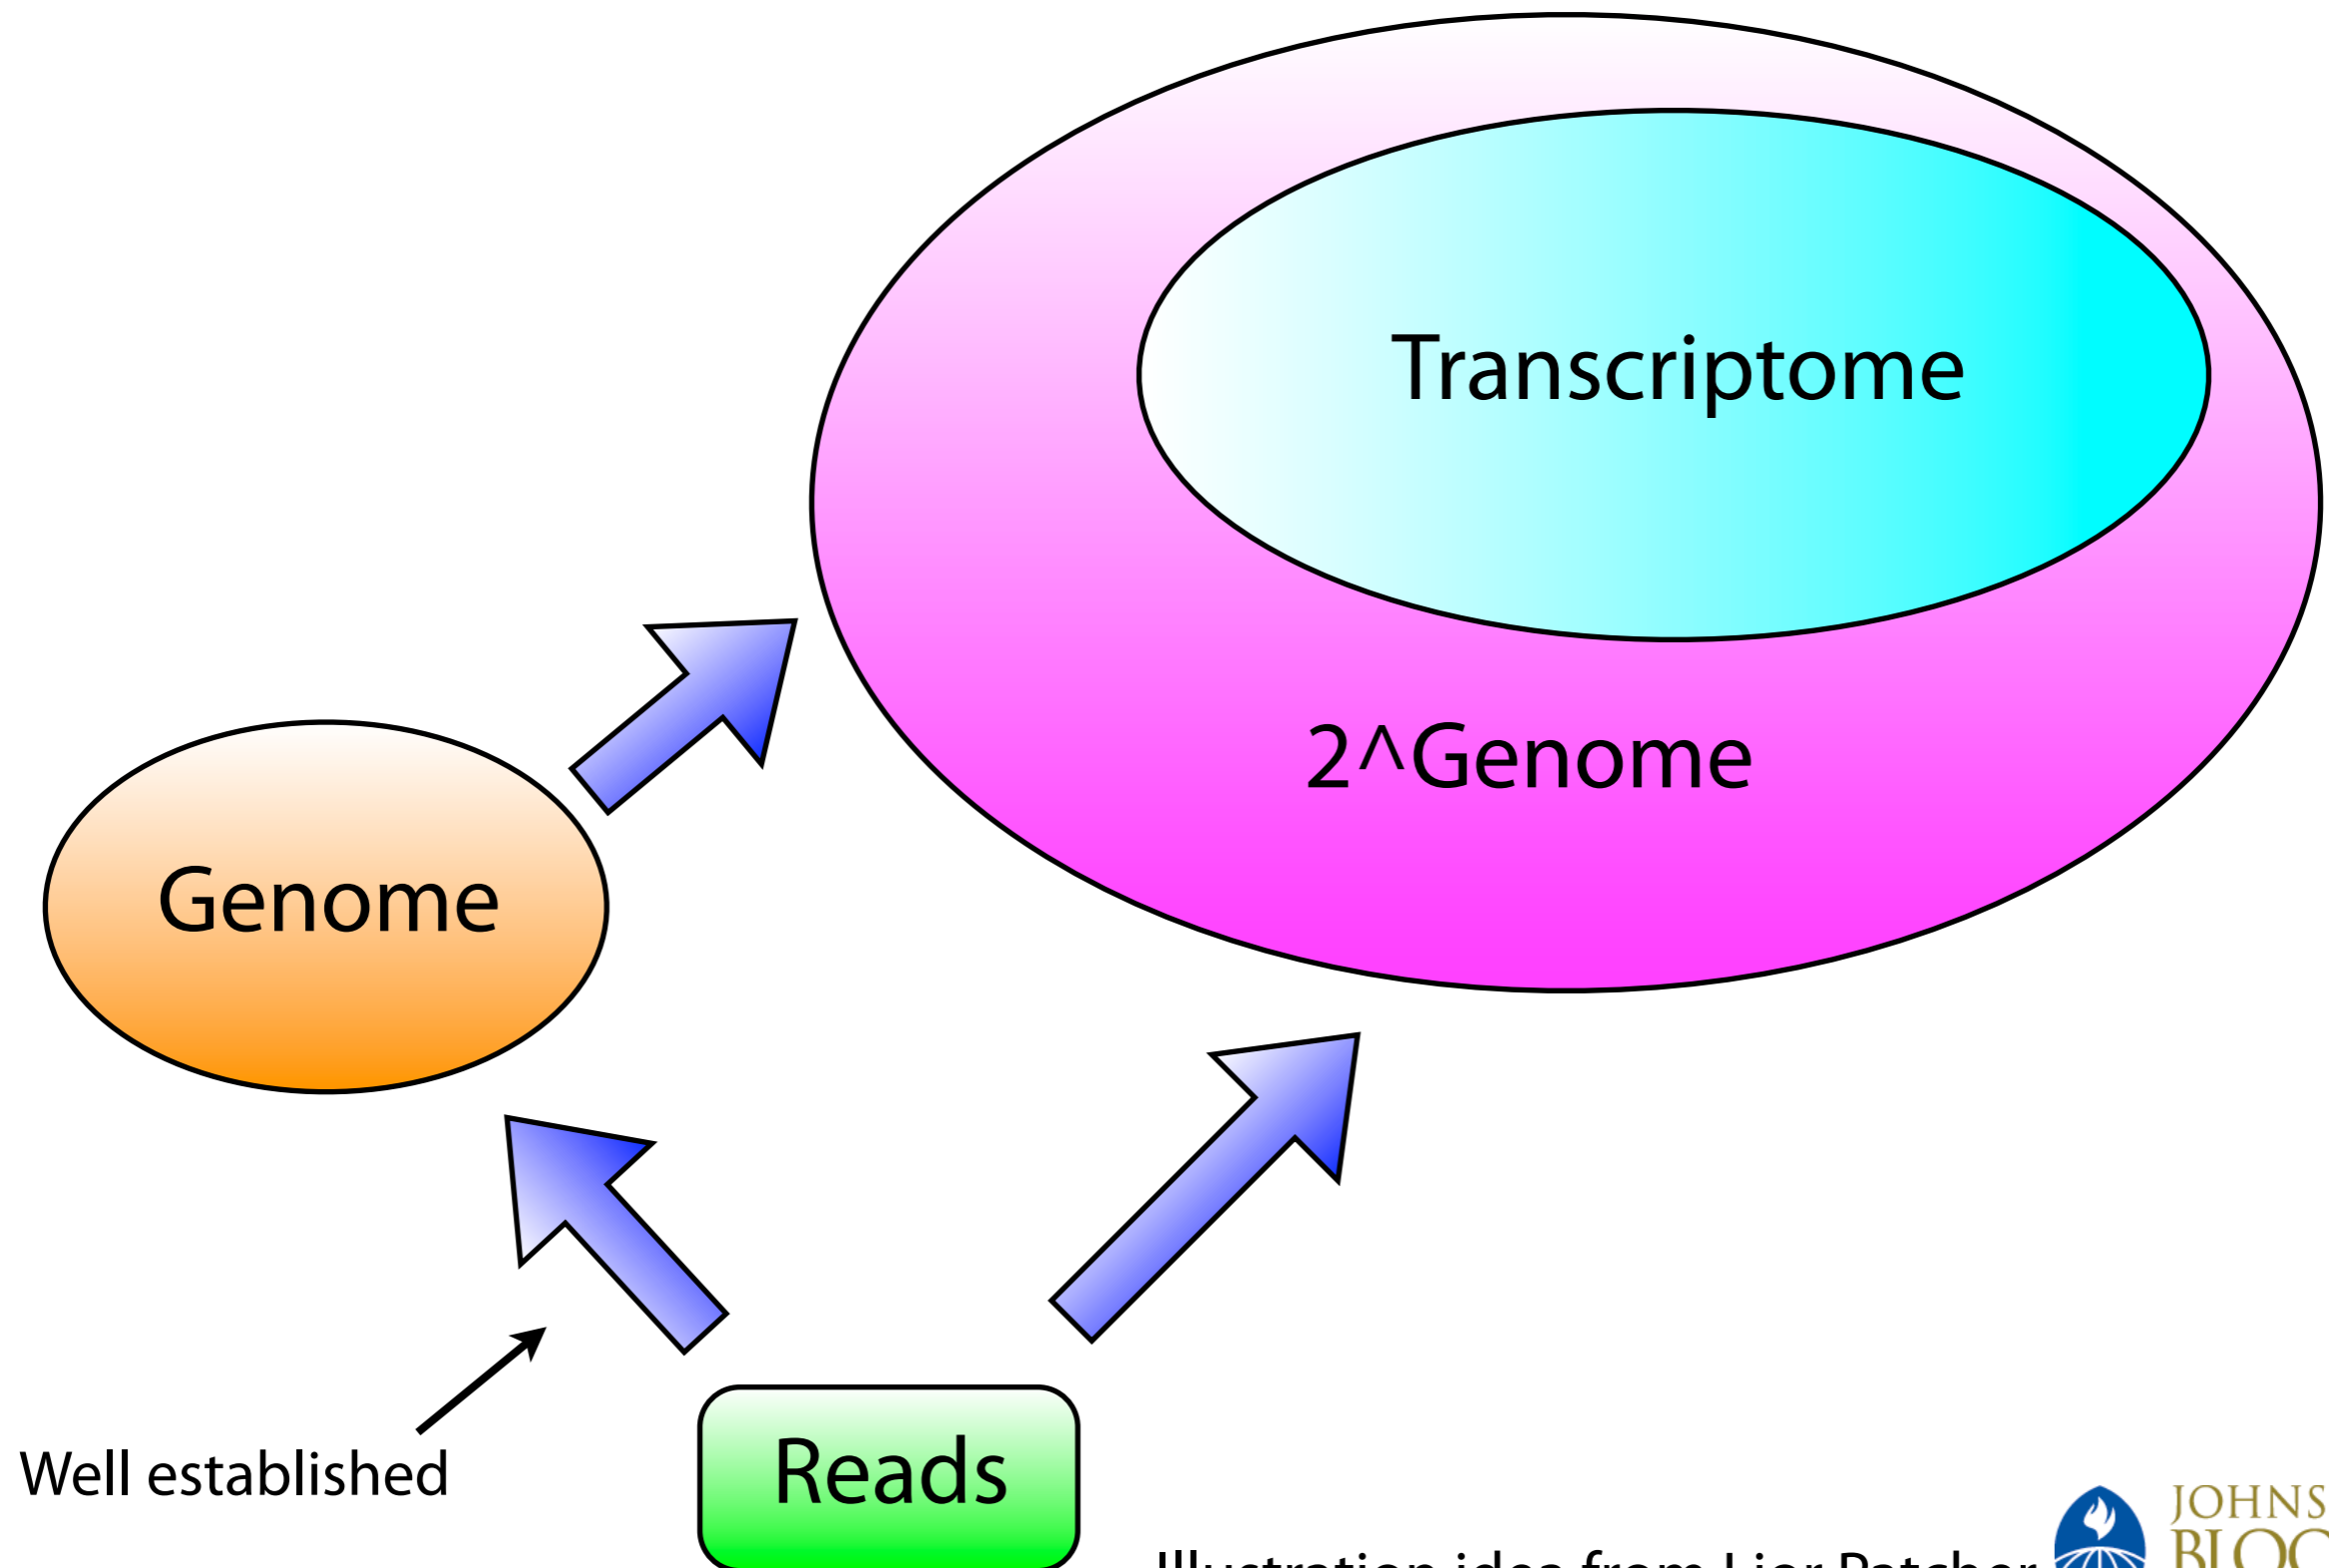


Illustration idea from Lior Patcher



The basic approaches


a

De novo assembly of the transcriptome




b

Map onto the genome



c

Map onto the genome and splice junctions



From Pepke (2009 Nat Methods)

Popular tools: Tophat/Cufflinks, GSNAP