


# Quality Assessment of sequencing data



- 
- Why QC?
    - Bad data = bad assembly
      - Partial / missing data?
      - Is there enough data?
      - Did I get what I expect?
      - Can I assemble it?
      - Do I need to change my analysis workflow?
      - Is it the correct type?
      - Are there biases?
      - Is there contamination?
    - Most checks can be made before assembly and assembly validation
  - Focus:
    - Illumina data (PE + MP)
    - Pacbio data
    - Nanopore data


- Make sure your data is whole.
  - File checksums ensure data integrity
    - MD5
      - `$ md5sum file1.fastq.gz # before`  
823fc8b0ca72c6e9bd8c5dcb0a66ce9b      file1.fastq.gz
      - `$ md5sum -c checksums.md5 # after`  
file1.fastq.gz: OK  
file2.fastq.gz: OK  
file3.fastq.gz: FAILED  
md5sum: WARNING: 1 of 3 computed checksums did NOT match
    - Calculate file checksums before transfer.
    - Validate checksums against the transferred files after the transfer.

- Paired end Illumina library






- Mate pair Illumina library





# Do I have enough data?

- What is my expected genome size?
- What depth of coverage should I expect?
  - Illumina:
    - 100x coverage in total
- Coverage = Number of bases sequenced / Estimated genome size



- FastQC / MultiQC summary reports
- Third party scripts
- Command line calculation (my favourite way)
  - Can use Seqtk to convert files to fasta
  - **zcat \*.fastq.gz | seqtk seq -A - | grep -v "^>" | tr -dc "ACGTNacgtn" | wc -m**
    - zcat ( concatenates the compressed fastq files into one stream )
    - seqtk ( converts to fasta format and drops reads less than 10k )
    - grep ( -v excludes lines starting with ">", i.e. fasta headers )
    - tr ( -dc removes any characters not in set "ACGTNacgtn" )
    - wc ( -m counts characters )
  - **parallel 'seqtk seq -A {} | grep -v "^>" | tr -dc "ACGTNacgtn" | wc -m' ::: \*.fastq.gz | paste -sd+ | bc -l**


- How much data is too much data?
  - Greater than 200X coverage is considered extreme.
- Why is too much data bad?
  - Increased computation time and resources
  - Errors begin to compound and start to look like real data.
  - Assemblies become more fragmented and inaccurate.
- How should I subsample?
  - Use a random fraction of the reads maintaining read pairing.
    - E.g. Use the same seed (-s) and give the fraction (0.1) in Seqtk.  
`seqtk sample -s100 read1.fq 0.1 > sub1.fq`  
`seqtk sample -s100 read2.fq 0.1 > sub2.fq`
  - Normalize uneven coverage (e.g. bbnorm)

- What does it tell you?
  - Total read pairs
  - Sequence length
  - Quality Score Encoding
  - Average GC%
  - Base quality along the read
  - Nucleotide % along the read
  - Sequence GC content
  - Duplication %
  - Adapter content
  
- Look at MultiQC for multiple samples


## Basic Statistics

Measure	Value
Filename	8361-F11_1.fastq.gz
File type	Conventional base calls
Encoding	Sanger / Illumina 1.9
Total Sequences	2809593
Sequences flagged as poor quality	0
Sequence length	300
%GC	39


## Per base sequence quality




## Per base sequence content



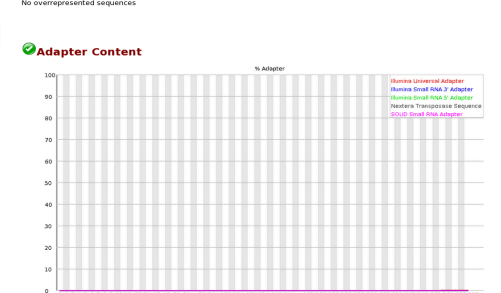
## Per sequence GC content



## Sequence Duplication Levels




## Overrepresented sequences



## Basic Statistics

Measure	Value
Filename	8361-F11_1.fastq.gz
File type	Conventional base calls
Encoding	Sanger / Illumina 1.9
Total Sequences	2809593
Sequences flagged as poor quality	0
Sequence length	300
%GC	39

## ✔ Per base sequence quality

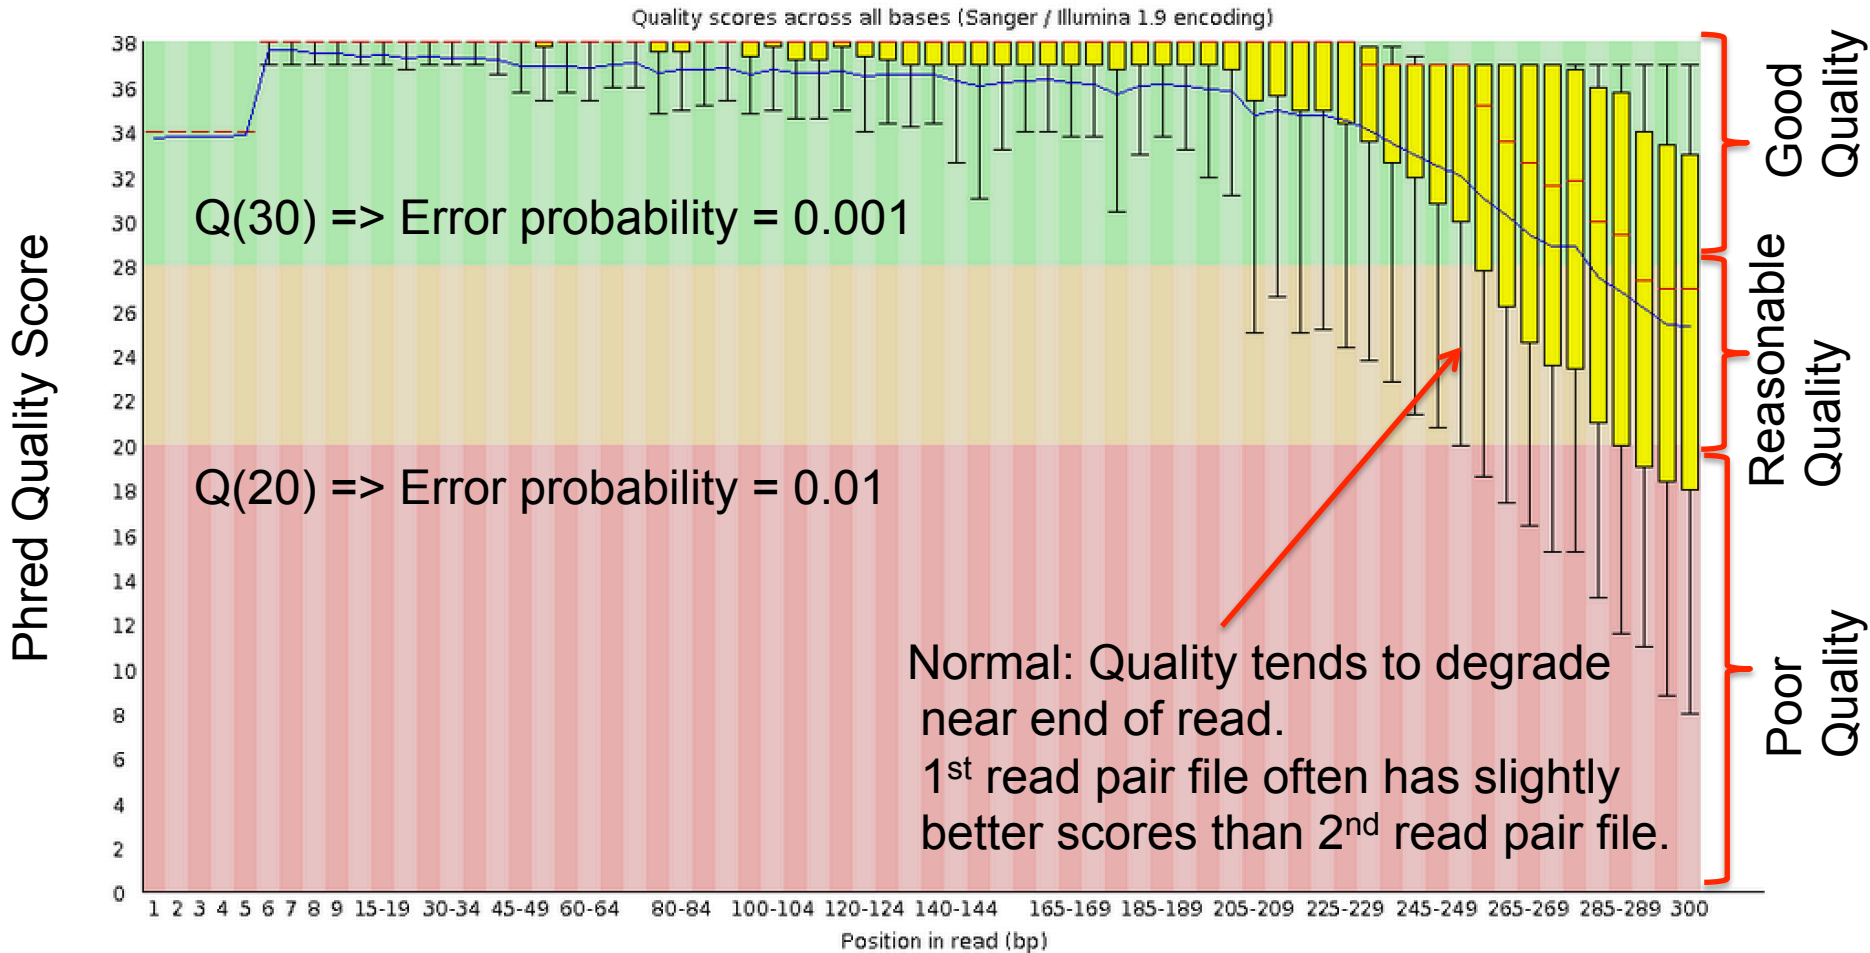





Q(40) => Error probability = 0.0001

## ✔ Per base sequence quality


Blue line = mean quality




## ✔ Per base sequence content



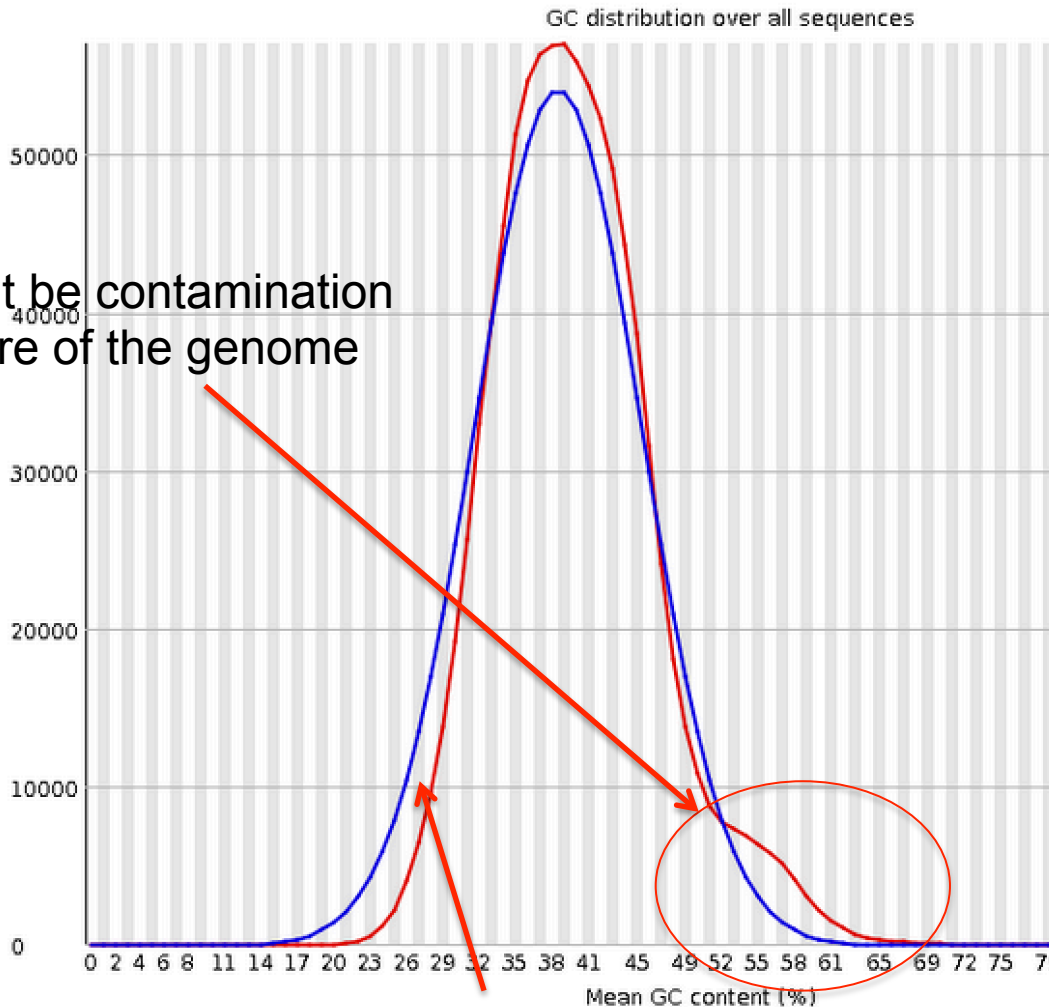
## ✔ Per base sequence content



## ⚠ Per sequence GC content




## ⚠ Per sequence GC content

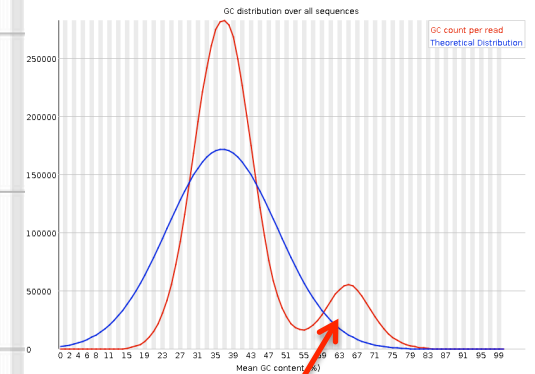


This might be contamination or a feature of the genome

Expected: Normal/Gaussian Distribution




Sharp peak indicates specific motif. Adapters are the usual suspect.



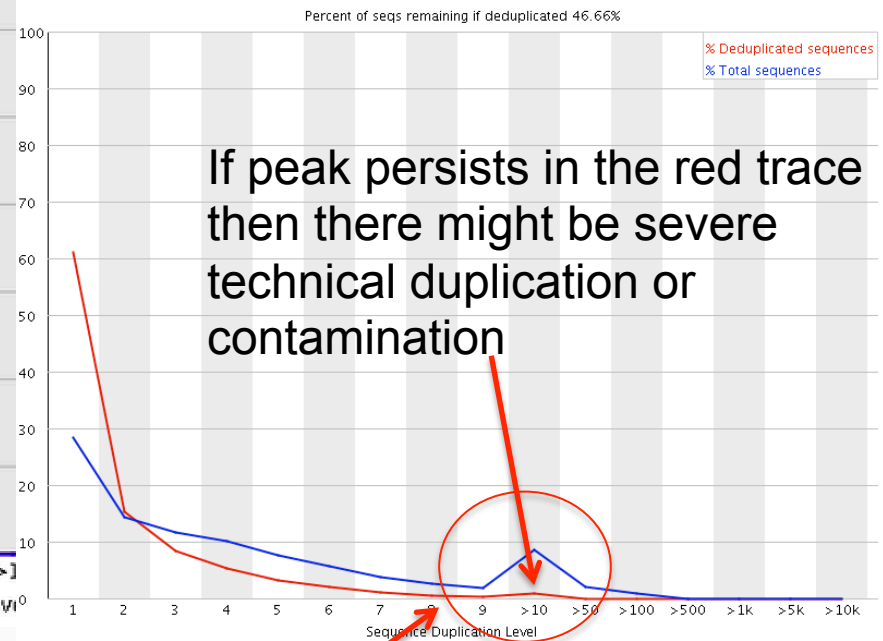
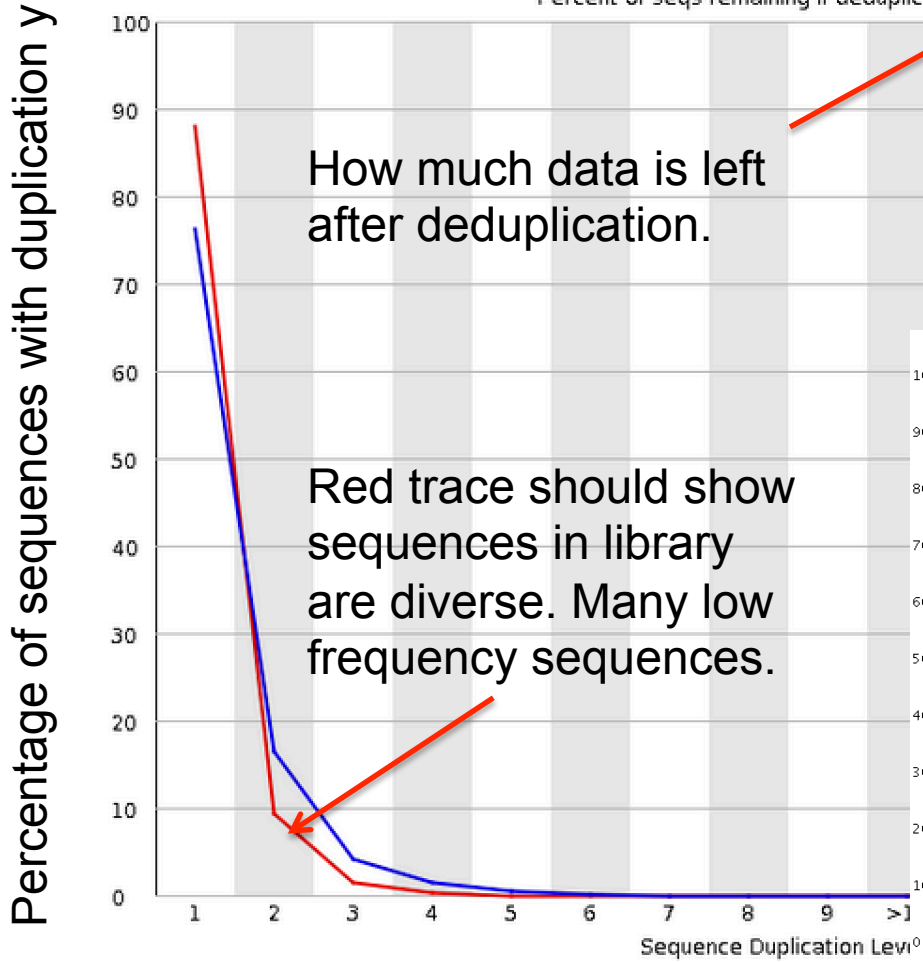
Wider or multiple distributions suggest contamination.

## ✔ Sequence Duplication Levels



## ✔ Sequence Duplication Levels

First 100,000 sequences tracked until end of file




Peak shows 10%+ sequences with high duplication levels

## ✔ Overrepresented sequences

No overrepresented sequences

## ✔ Adapter Content





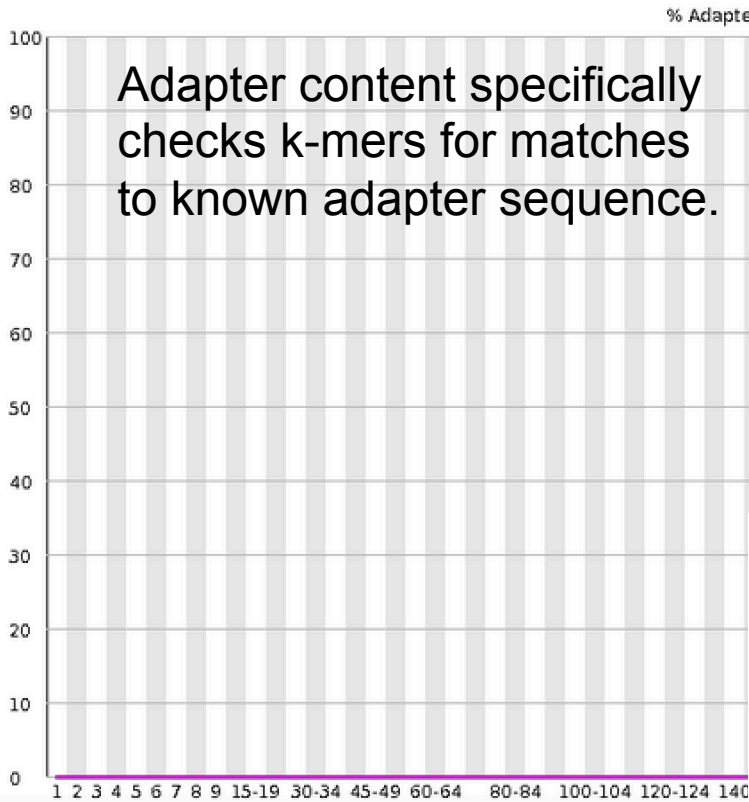
## ✔ Overrepresented sequences

No overrepresented sequences

Lists sequence that is more than 0.1%

First 100,000 sequences tracked until end of file

## ✔ Adapter Content



Adapter content specifically checks k-mers for matches to known adapter sequence.

Overrepresented sequences are matched against known contaminants.


Match hits are not conclusive, but indicative.

Matches must be >20bp and only 1 mismatch.

## ✘ Overrepresented sequences

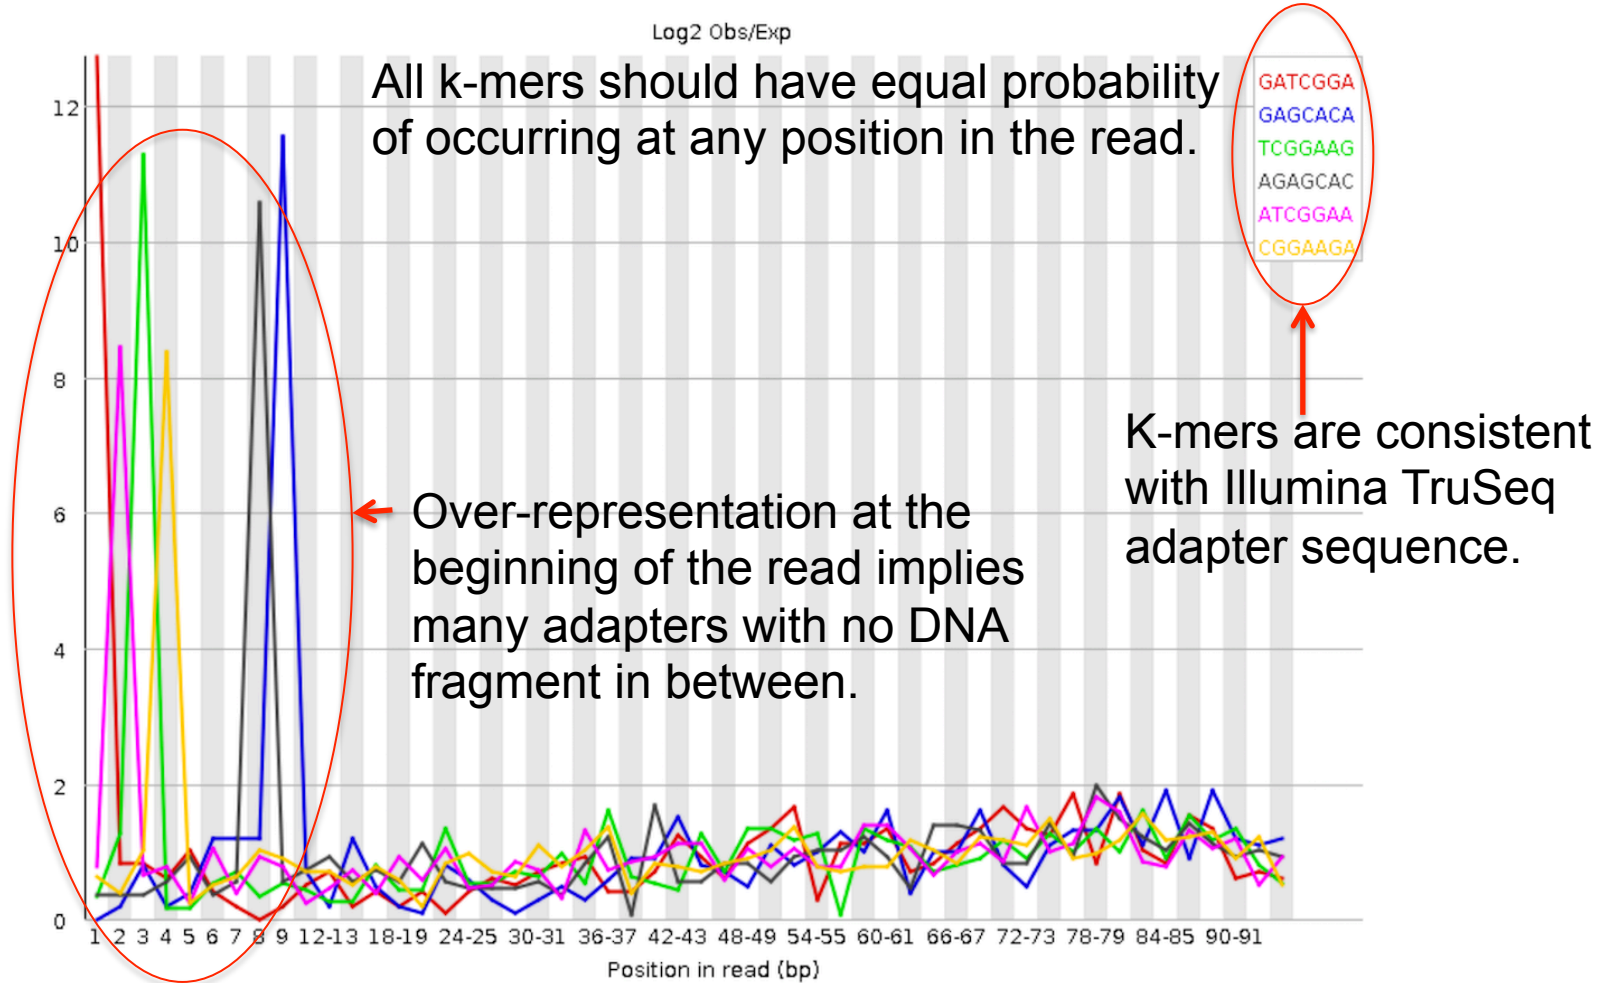
Sequence	Count	Percentage	Possible Source
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGACCGATCT	8122	8.122	Illumina Paired End PCR Primer 2 (100% over 40bp)
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGATCGGAAG	5086	5.086	Illumina Paired End PCR Primer 2 (97% over 36bp)
AATGATACGGCGACCACCGAGATCTACACTCTTTCCCTAC	1085	1.085	Illumina Single End PCR Primer 1 (100% over 40bp)
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGACCGGAAG	508	0.508	Illumina Paired End PCR Primer 2 (97% over 36bp)
AATTATACGGCGACCACCGAGATCTACACTCTTTCCCTAC	242	0.242	Illumina Single End PCR Primer 1 (97% over 40bp)
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGATCGGAA	235	0.23500000000000001	Illumina Paired End Adapter 2 (96% over 31bp)
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGATCGGAAG	228	0.227999999999999998	Illumina Paired End Adapter 2 (96% over 28bp)
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGACCGGAGC	205	0.205000000000000002	Illumina Paired End PCR Primer 2 (97% over 36bp)
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGATCGGAA	183	0.183	Illumina Paired End Adapter 2 (100% over 32bp)
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGATCGGAAG	183	0.183	Illumina Paired End Adapter 2 (100% over 32bp)
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGACCGAATCT	164	0.164	Illumina Paired End PCR Primer 2 (97% over 40bp)
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGACCGGTCT	129	0.129	Illumina Paired End PCR Primer 2 (97% over 40bp)
AATTATACCTTACCACCTATATCTACACTCTTTCCCTAC	123	0.123	No Hit
GATCGGAAGAGCGGTTTCAGCAGGAATGCCGAGACCGGACT	122	0.122	Illumina Paired End PCR Primer 2 (97% over 36bp)
CGGTTTCAGCAGGAATGCCGAGATCGGAAGAGCGGTTTCAGC	113	0.112999999999999999	Illumina Paired End PCR Primer 2 (96% over 25bp)

## ✖ Kmer Content




## ✖ Kmer Content

Is a k-mer over-represented along the length of a read?




Default k is 7. K-mer size can be increased with option -k





# Trimming reads

- Why trim reads?
  - Remove poor quality reads




- Many tools available
  - Trimmomatic
  - CutAdapt
  - AlienTrimmer
  - Sickle
  - Trim Galore
  - Scythe
  - Prinseq
  - ...
- **Warning:** Some assemblers expect untrimmed input
  - Allpaths-LG
  - Mira

- Trimmomatic:  
java -jar trimmomatic-0.36.jar PE -phred33  
input\_forward.fq.gz input\_reverse.fq.gz  
output\_forward\_paired.fq.gz  
output\_forward\_unpaired.fq.gz  
output\_reverse\_paired.fq.gz  
output\_reverse\_unpaired.fq.gz  
ILLUMINACLIP:TruSeq3-PE.fa:2:30:10 LEADING:3  
TRAILING:3 SLIDINGWINDOW:4:15 MINLEN:36
- Library type
- Quality Score encoding
- Input file pair
- Output
- Fasta file of sequences to remove (adapters, linkers, etc)
- How to trim and what to keep

- BBMerge can be used to discover adapters:  
bbmerge.sh in=reads.fq outa=adapters.fa



- Why do duplicates arise?
  - Optical duplicates (amplified cluster mistaken for multiple clusters)
  - PCR duplicates
- Why are duplicates bad?
  - Poor overlap information
  - Increased variance of coverage
  - Increased computation time and resources
- How to remove duplicates:
  - Prinseq
  - FastUniq
  - ParDRe
  - ...



SMRTbell™ Template



## Polymerase Read

### Definition:

- Sequence of nucleotides incorporated by polymerase while reading a template
- Includes adapters
- Often called “read”
- Includes adapters
- 1 molecule, 1 pol. read

### Purpose:

- QC of instrument run
- Benchmarking



## Subread

### Definition:

- Single pass of template
- Adapters removed
- 1 molecule,  $\geq 1$  subreads

### Unique data:

- Kinetic measurements
- Rich QVs

### Purpose:

- For subsequent analysis



## Read of Insert

### Definition:

- Represents highest-quality single-sequence for an insert, regardless of number of passes
- Generalizes CCS for  $<2$  passes and  $RQ < 0.9$
- 1 or more passes
- 1 molecule, 1 read

### Purpose:

- For Library QC
- For subsequent analysis

```
m140415_143853_42175_c100635972550000001823121909121417_s1_p0/553/3100_11230
└─1┤ └─2──┤ └─3┤ └──────────────────────────4──────────────────────────┤ └─5┤ └─6┤ └─7┤ └─8──┤
```

1. " m " = movie
2. Time of Run Start ( `yymmdd_hhmmss` )
3. Instrument Serial Number
4. SMRT Cell Barcode
5. Set Number (a.k.a. "Look Number". Deprecated field, used in earlier version of RS)
6. Part Number (usually " `p0` ", " `x0` " when using expired reagents)
7. ZMW hole number †
8. Subread Region ( `start_stop` using polymerase read coordinates) †

† Note that Fields 7 and 8 are used as sequence IDs in FASTA|FASTQ files. They are not used in filenames.

```
@m150619_093250_42174_c100795682550000001823166309091510_s1_p0/109/0_4936 RQ=0.879
@m150619_093250_42174_c100795682550000001823166309091510_s1_p0/109/4981_9942 RQ=0.879
@m150619_093250_42174_c100795682550000001823166309091510_s1_p0/109/9988_10378 RQ=0.879
@m150619_093250_42174_c100795682550000001823166309091510_s1_p0/157/0_7588 RQ=0.871
@m150619_093250_42174_c100795682550000001823166309091510_s1_p0/157/7628_15139 RQ=0.871
@m150619_093250_42174_c100795682550000001823166309091510_s1_p0/157/15186_22778 RQ=0.871
@m150619_093250_42174_c100795682550000001823166309091510_s1_p0/157/22820_30464 RQ=0.871
@m150619_093250_42174_c100795682550000001823166309091510_s1_p0/157/30510_36641 RQ=0.871
```

- The subreads fastq file contains all the subreads from a SMRT movie.
- The reads from a ZMW after adapter removal are oriented in the direction forward, reverse, forward, and so on.
- Read Quality (RQ) Assignment: A trained prediction of a read's mapped accuracy based on its pulse and base file characteristics (peak signal-to-noise ratio, average base QV, interpulse distance, and so on).
- Quality Value (QV): The total probability that the basecall is an insertion or substitution or is preceded by a deletion.  $QV = -10 * \log_{10}(p)$ .

# Do I have enough data?

- What is my expected genome size?
- What depth of coverage should I expect?
  - PacBio:
    - 70x coverage in total from subreads per allele
    - At least 30x coverage of reads >10kb per allele
- Coverage = Number of bases sequenced / estimated genome size

11/30/2015

Reports for Job pb\_251\_1\_subreads\_CTR

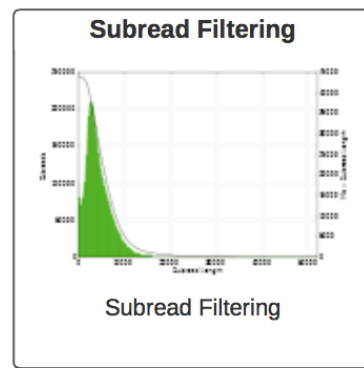
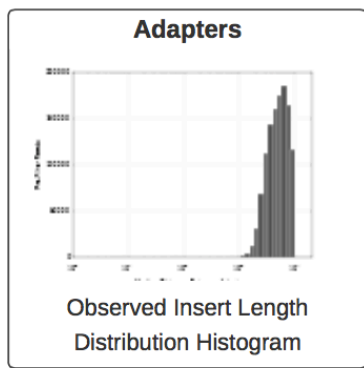
## Reports for Job pb\_251\_1\_subreads\_CTR



SMRT Cells: 72 Movies: 72



### Overview

Job Metric	Value
Adapter Dimers (0-10bp)	0.06%
Short Inserts (11-100bp)	0.01%
Number of Bases	44,946,763,242
Number of Reads	3,918,307
N50 Read Length	24,367
Mean Read Length	11,470
Mean Read Score	0.85





- Third party scripts
- Command line calculation (my favourite way)
  - Can use Seqtk to convert and filter on read length
  - **zcat \*.fastq.gz | seqtk seq -A -L 10000 - | grep -v "^>" | tr -dc "ACGTNacgtn" | wc -m**
    - zcat ( concatenates the compressed fastq files into one stream )
    - seqtk ( converts to fasta format and drops reads less than 10k )
    - grep ( -v excludes lines starting with ">", i.e. fasta headers )
    - tr ( -dc removes any characters not in set "ACGTNacgtn" )
    - wc ( -m counts characters )
  - **parallel 'seqtk seq -A -L 10000 {} | grep -v "^>" | tr -dc "ACGTNacgtn" | wc -m' ::: \*.fastq.gz | paste -sd+ | bc -l**


Job Metric	Value
Adapter Dimers (0-10bp)	0.06%
Short Inserts (11-100bp)	0.01%
Number of Bases	44,946,763,242
Number of Reads	3,918,307
N50 Read Length	24,367
Mean Read Length	11,470
Mean Read Score	0.85



**Filtering**

Metrics	Pre-Filter	Post-Filter
Polymerase Read Bases	49236076578	44946763242
Polymerase Reads	10821024	3918307
Polymerase Read N50	23758	24367
Polymerase Read Length	4550	11470
Polymerase Read Quality	0.319	0.846





#### Adapters

Adapter Dimers (0-10bp)	0.06%
Short Inserts (11-100bp)	0.01%



## Loading

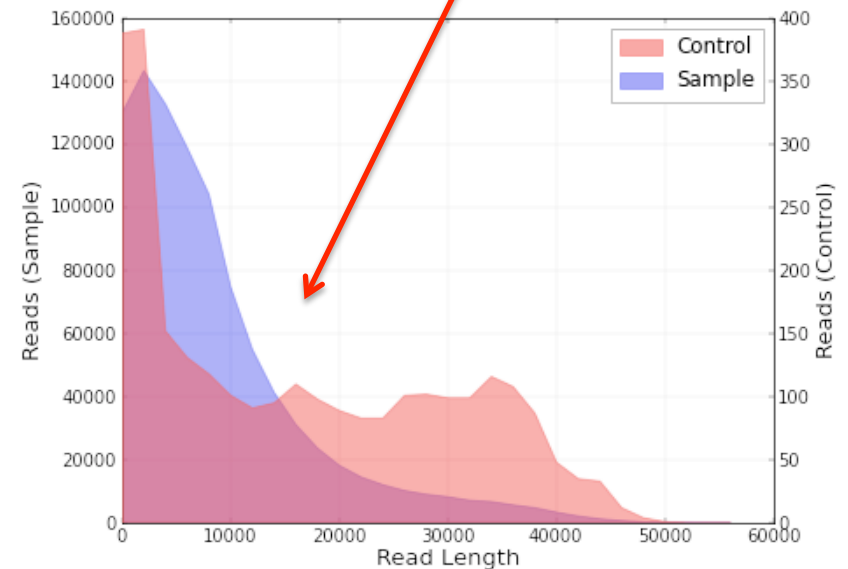
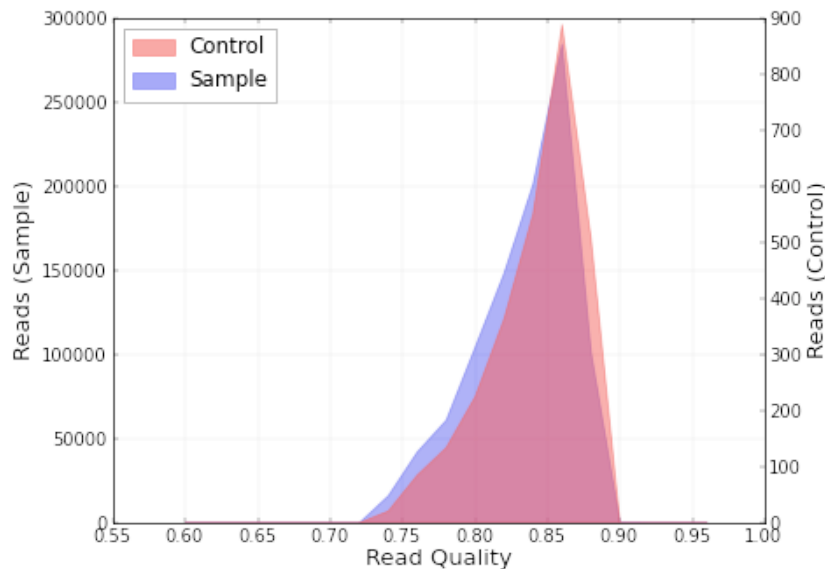
SMRT Cell ID	Productive ZMWs	ZMW Loading For Productivity 0	ZMW Loading For Productivity 1	ZMW Loading For Productivity 2
m151122_235521_42203_c100927002550000001823210705121641	150,292	50.73%	40.19%	9.08%
m151124_195105_42237_c100966232550000001823205304301611	150,292	40.75%	51.31%	7.94%
m151122_151707_42203_c100927102550000001823210705121617	150,292	57.69%	33.55%	8.75%
m151114_001837_42237_c100926912550000001823210705121673	150,292	56.6%	31.53%	11.87%
m151105_141536_42237_c100884702550000001823198604021655	150,292	35.48%	55.12%	9.4%
m151107_172533_42237_c100926842550000001823210705121675	150,292	40.2%	46.18%	13.63%
m151123_082023_42237_c100927112550000001823210705121606	150,292	61.16%	31.51%	7.34%
m151125_042931_42237_c100966232550000001823205304301613	150,292	44.14%	47.93%	7.93%

- SMRT cell loading
  - P0: % of ZMWs that are empty with no polymerase
  - P1: % of ZMWs that are productive and sequencing
  - P2: % of ZMWs that are not P0 or P1 (e.g. unbound polymerase, more than one molecule in a well (overloaded cell)).
  - Maximize P1 and minimize P0 + P2.
  - High P0 indicates underloading (too low concentration of molecules)
  - High P2 indicates overloading (too high concentration) or poor prep.

- SMRT portal does not filter out control reads unless the protocol is included.

Control Sequence	2kb_Control	Number of Control Reads	2775
Fraction Control Reads	0.0028828862522167057	Control Subread Accuracy	0.861804283567377
Control Polymerase Read Length N50	29699.0	Control Polymerase Read Length 95%	39372
Control Polymerase Read Length	14950		


Control reads are longer than Sample reads indicating good sequencing but bad Sample DNA quality




# Adapter Trimming

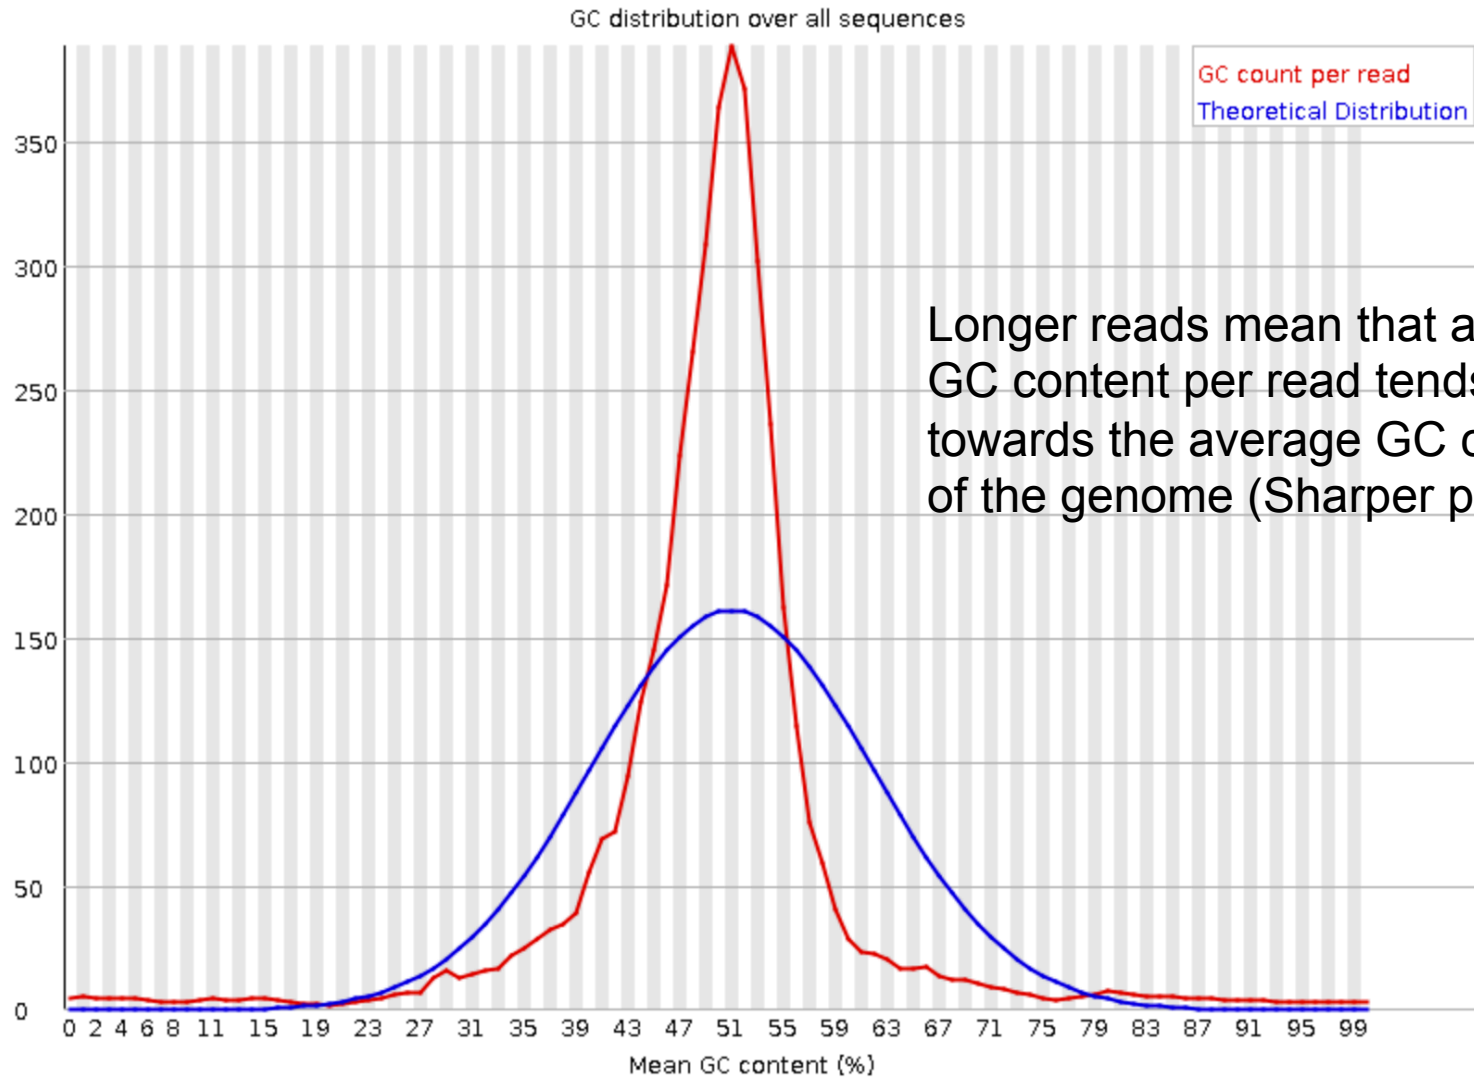
SMRTbell adapter:

ATCTCTCTCTTTTCCTCCTCCTCCGTTGTTGTTGTTGAGAGAGAT



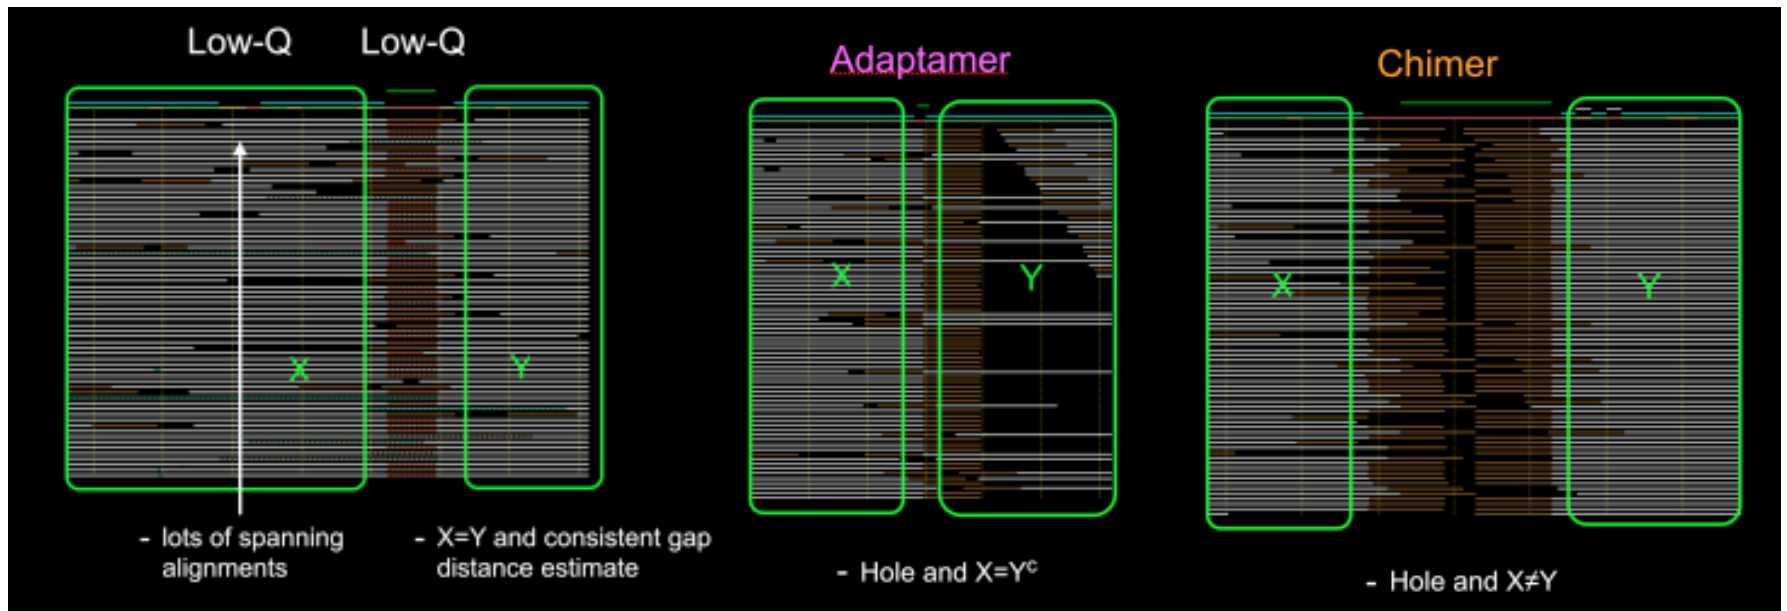

Quality scores across all bases (Sanger / Illumina 1.9 encoding)






# Read Errors

Quality score is inferred from the light signal.  
Strong light signal does not imply correct base calls.



# Nanopore sequencing

- Single Molecule sequencing
- No fragmentation required, therefore read length is theoretically unlimited. (Pacbio has size selected library)
- Support is limited to the ONT community portal.
- Limited QC tools.
  - Poretools (output summary)
  - Porechop (adapters)



- 
- Illumina
    - Md5sum
    - Format
    - Data Quantity
    - Quality Scores
    - GC Content
    - Nucleotide bias
    - Duplication
    - Adapter content
    - Fragment distribution
  - PacBio / Nanopore
    - Md5sum
    - Format
    - Data Quantity
    - GC Content
    - Adapter & Control check
    - Subread distribution



- 
- Sequencing Fail  
<https://sequencing.qcfail.com>
  - SEQanswers  
<http://seqanswers.com>
  - BioStar  
<https://www.biostars.org>
  - Bioinformatics StackExchange  
<https://bioinformatics.stackexchange.com>