

Bi5444 Analysis of sequencing data

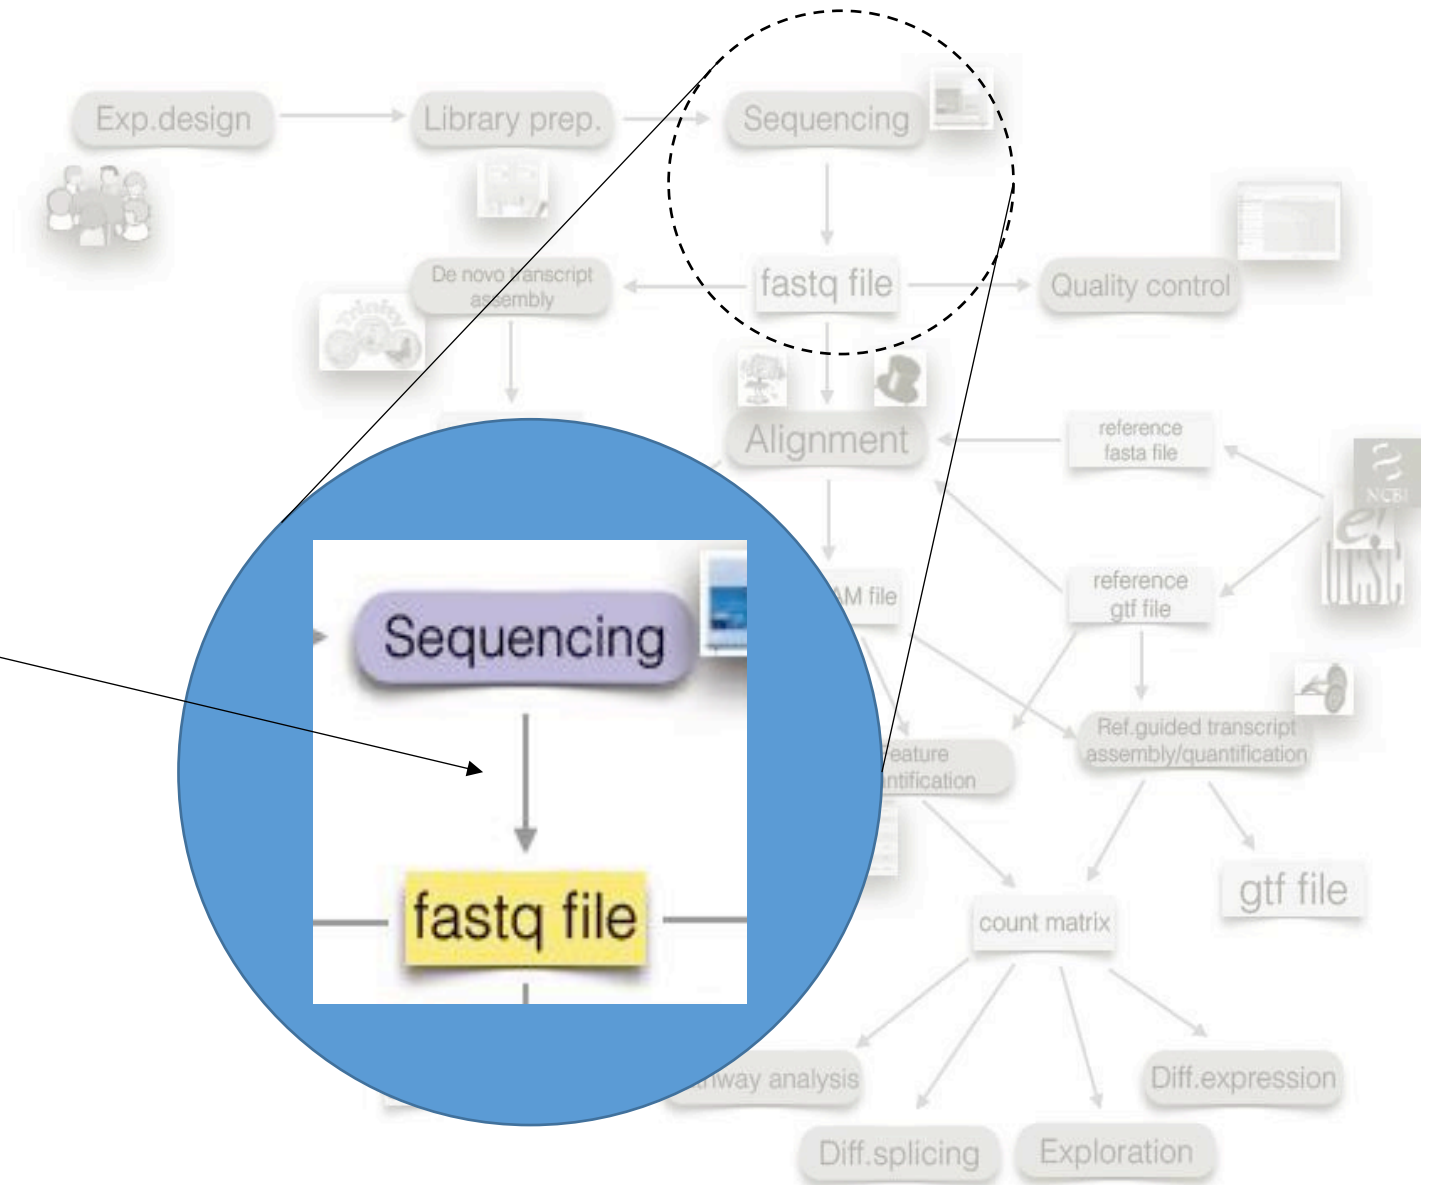
NGS quality control

---

Eva Budinska




Step 0:  
base calling  
(image analysis)  
+ base quality  
control



## Step 0: base calling (image analysis)

- The identity of each base of a cluster is read off from **sequential images**
- One cycle -> one image



# The PHRED score

$$Q_{phred} = - 10 \times \log_{10}P(\text{error})$$

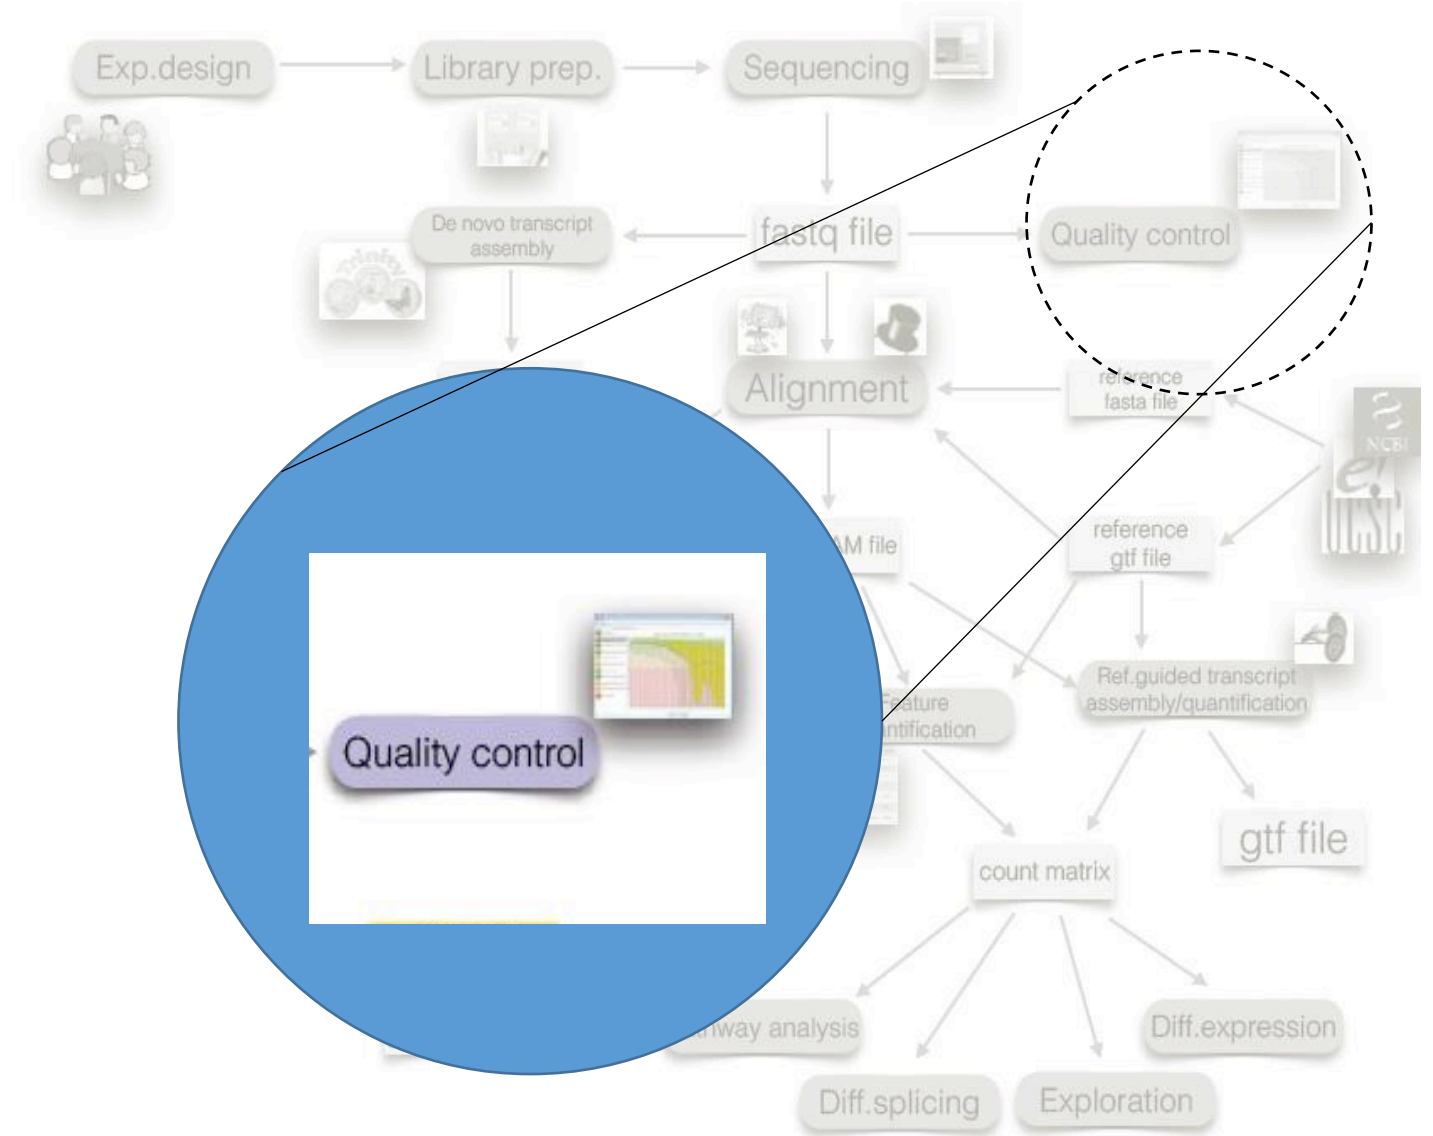
- The *Phred* quality score is the negative ratio of the error probability to the reference level of  $P = 1$  expressed in Decibel (dB).
- The **error estimate** is based on **statistical model** providing measure of **certainty** of each base call in addition to the nucleotide itself
- These statistical models base their error estimate on:
  - Signal intensities from the recorded image
  - Number of the sequencing cycle
  - Distance to other sequence colonies
- *Phred* score is recoded using ASCII in fastq file

Phred score	Probability of incorrect base call	Base call
10	1 in 10	90%
20	1 in 100	99%
30	1 in 1000	99.9%
40	1 in 10 000	99.99%
50	1 in 100 000	99.999%
60	1 in 1 000 000	99.9999%





# Step 1: Read quality control and data filtering






# Before we dive in...

... let's review few concepts and expressions


# The steps of Illumina sequencing

1. Fragment genomic DNA, e.g. with a sonicator.
2. Ligate adapters to both ends of the fragments.
3. PCR amplify the fragments with adapters
4. Spread DNA molecules across flowcells. Goal is to get exactly **one DNA molecule** per flowcell lawn of primers. This depends purely on probability, based on the concentration of DNA.
5. Use bridge PCR to amplify the single molecule on each lawn so that you can get a strong enough signal to detect. Usually this requires several hundred or low thousands of molecules.
6. Sequence by synthesis of complementary strand: [reversible terminator chemistry](#).


## A. Library Preparation



## A. Cluster Amplification




## C. Sequencing



Sequencing reagents, including fluorescently labeled nucleotides, are added and the first base is incorporated. The flow cell is imaged and the emission from each cluster is recorded. The emission wavelength and intensity are used to identify the base. This cycle is repeated "n" times to create a read length of "n" bases.


## D. Alignment & Data Analysis



Reads are aligned to a reference sequence with bioinformatics software. After alignment, differences between the reference genome and the newly sequenced reads can be identified.

# Sources of errors: adapters

- In step 2, adapters are ligated to the end of the fragments




Sequencing random fragments of DNA is possible via the addition of short nucleotide sequences which allow any DNA fragment to:

- Bind to a flow cell for next generation sequencing
- Allow for PCR enrichment of adapter ligated DNA fragments only
- Allow for indexing or 'barcoding' of samples so multiple DNA libraries can be mixed together into 1 sequencing lane (known as multiplexing)

# Sources of errors: PCR duplicates

- In step 3 we are *intentionally* creating multiple copies of each original genomic DNA molecule so that we have enough of them.
- PCR duplicates occur when **two copies of the same original molecule get onto different primer lawns in a flowcell.**
- In consequence we read the very same sequence twice!

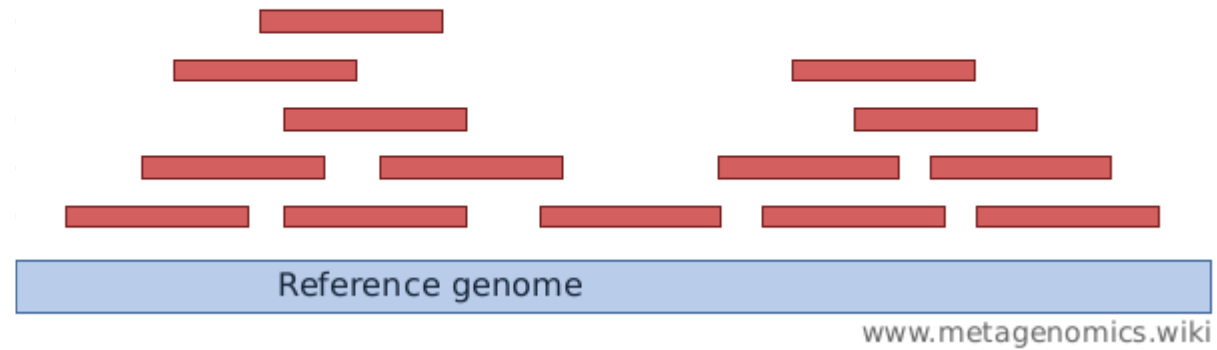
Higher rates of PCR duplicates e.g. 30% arise when you have too little starting material such that greater amplification of the library is needed in step 3, or when you have too great a variance in fragment size, such that smaller fragments, which are easier to PCR amplify, end up over-represented.





# Sequencing coverage

**Coverage** in DNA sequencing is the number of unique reads that include a given nucleotide in the reconstructed sequence.



# Depth of coverage

(coverage depth / mapping depth)

*How strongly is the genome "covered" by sequenced fragments (short reads)?*

**Per-base coverage** is the average number of times a base of a genome is sequenced (in other words, how many reads cover it).

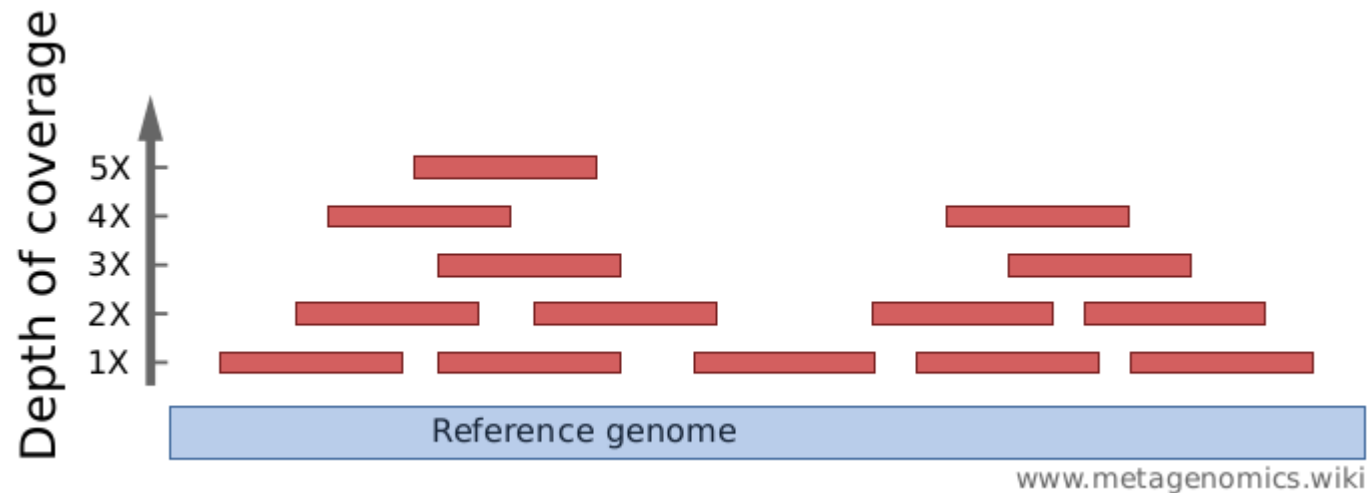
Average coverage of the genome ( $A_v$ )

$$A_v = (N \times L) / G$$

$G$  - length of the original genome

$N$  - number of reads


$L$  - average read length



**The coverage depth of a genome** is calculated as **the number of bases of all short reads that match a genome divided by the length of this genome**. It is often expressed as 1X, 2X, 3X,... (1, 2, or 3 times coverage).

## Breadth of coverage (covered length)

*What proportion of the genome is "covered" by short reads?  
Are there regions that are not covered, even not by a single  
read?*



**Breadth of coverage is the percentage of bases of a reference genome that are covered with a certain depth.**  
For example: "90% of a genome is covered at 1X depth; and still 70% is covered at 5X depth."





# Sequencing coverage

- **Deep sequencing** refers to the general concept of aiming for high number of unique reads of each region of a sequence.

## Step 1: Read quality control and data filtering

- Uses the output file with information about the quality of base calls (.fastq)
- First step in the pipeline that **deals with actual sequencing data** in base or color space

- Several metrics are evaluated, not all of them use the Phred score information:
  - Distribution of quality scores at each sequence, Sequence composition, Per-sequence and per-read distribution of GC content, Library complexity, Overrepresented sequences
- Initial overview – already in base calling SW
- More quality overview – SW solutions `SolexaQA`, `FastQC`

# Main quality control points

1. **Base quality**
2. **Sequence composition** – sequence content across bases should not change with cycle (exception are targeted sequencing SNP experiments)
3. **Per-sequence and per-read distribution of GC content** (shift from expected can indicate contamination by rRNA for instance)
4. **Library complexity** (too many duplicates?)
5. **Overrepresented sequences** – may represent highly expressed genes, or presence of adapters or rRNA contamination or PCR duplicates



# Base quality


---

- Quality of bases (Phred score) should be good across all cycles  
(all the sequence)

# Base quality – an excellent example

---


- Shows distribution (boxplot) of quality of bases (Phred scores) across all reads in each cycle



# Base quality – a more common example

---


- Decrease of quality towards the end of reads (late cycles)



# Base quality – bad example

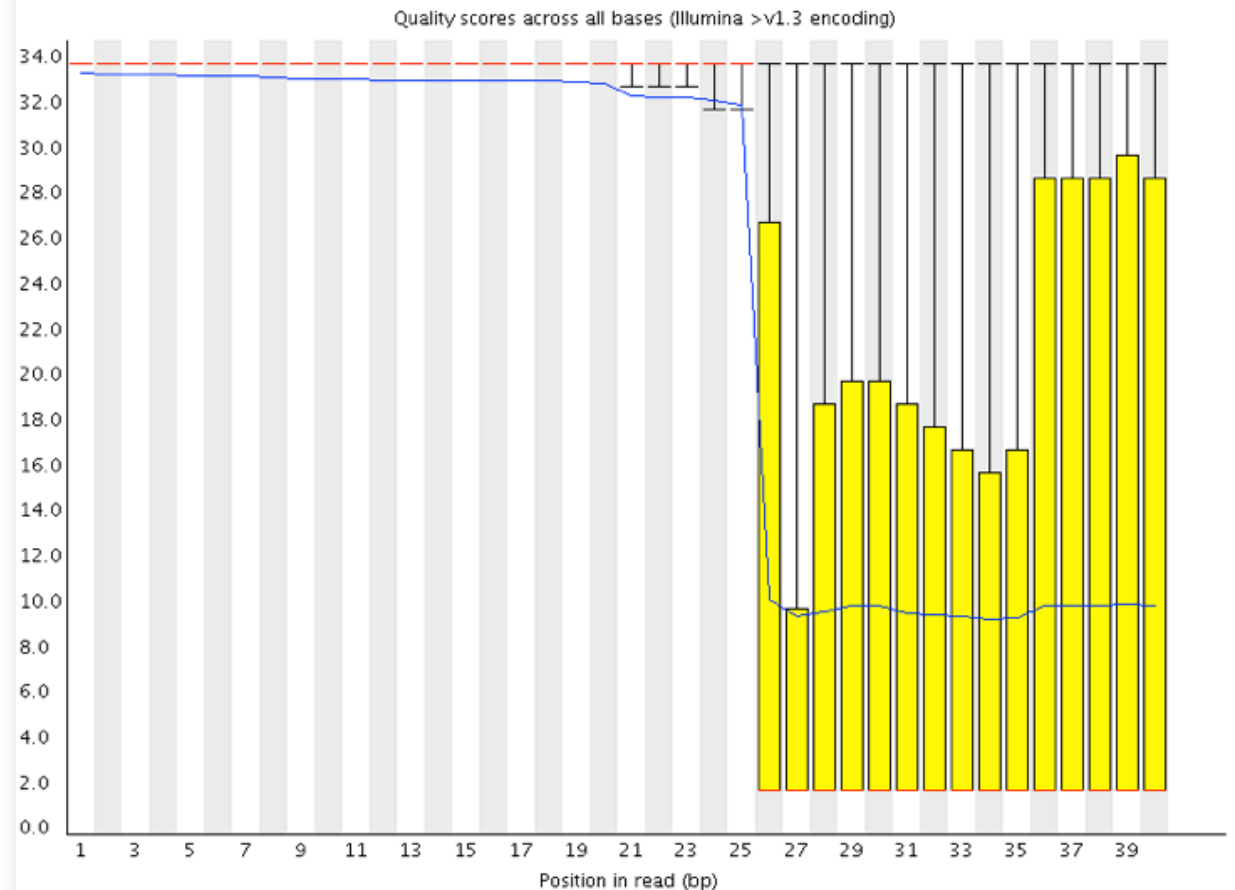
---

- ...



# Base quality - sudden quality drop


- Indicates problems with flow cell, trimming needed






# Base quality – targeted sequencing

- The low quality extremes suggest a problem in the beginning of the reads
- (primers?, NNNNN sequences...)



# Base quality — microbiome

- This is what it can look like with very small sample and sequence size



# Base call errors in last cycles

- Towards the end of sequencing, the quality drops, signal is worse
- We can see it for Illumina and SOLID
- Not very important for RNAseq, but crucial for variant calling

# SNP calling

TTAACAGTGTTTCAGTAAGTT  
TTCAGTAAGATTCCATGAGCT  
ACAGTGTTTCAGTAAGATTCC  
TGTTTCAGTAAGATTCCATGAG  
CAGTGTTTCAGTAAGTTTCCA  
GTTTCAGTAAGATTCCATGAGC  
GTGTTTCAGTAAGATTCCATGA  
TCAGTAAGATTCCATGAGCTC  
AACAGTGTTTCAGTAAGTTT  
AGTAAGATTCCATGAGCTCT  
CTTAACAGTGTTTCAGTAAGATTCCATGAGCTCT

---

# SNP calling

TTAACAGTGTTTCAGTAAGTT  
AACAGTGTTTCAGTAAGTTT  
ACAGTGTTTCAGTAAGATTCC  
CAGTGTTTCAGTAAGTTTCCA  
GTGTTTCAGTAAGATTCCATGA  
TGTTTCAGTAAGATTCCATGAG  
GTTTCAGTAAGATTCCATGAGC  
TTCAGTAAGATTCCATGAGCT  
TCAGTAAGATTCCATGAGCTC  
AGTAAGATTCCATGAGCTCT

---


CTTAACAGTGTTTCAGTAAGATTCCATGAGCTCT

# SNP error dependent on cycle

These errors are not random and look like SNPs (e.g. if there were randomly distributed T, C, G and A's, we would conclude it is error directly)

We want the SNPs to be distributed evenly across cycles

SNPs coming from towards end of the read are sign of false positive

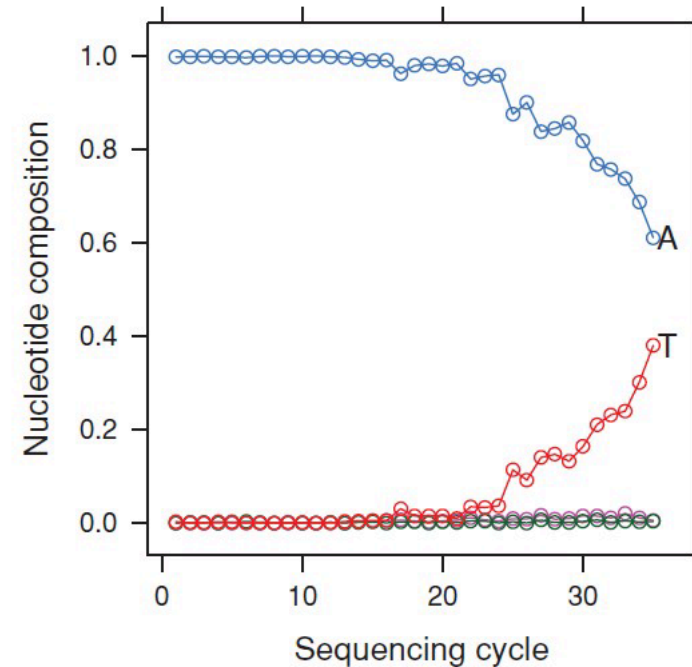


# SNP error dependent on cycle

These errors are not random and look like SNPs (e.g. if there were randomly distributed T, C, G and A's, we would conclude it is error directly)

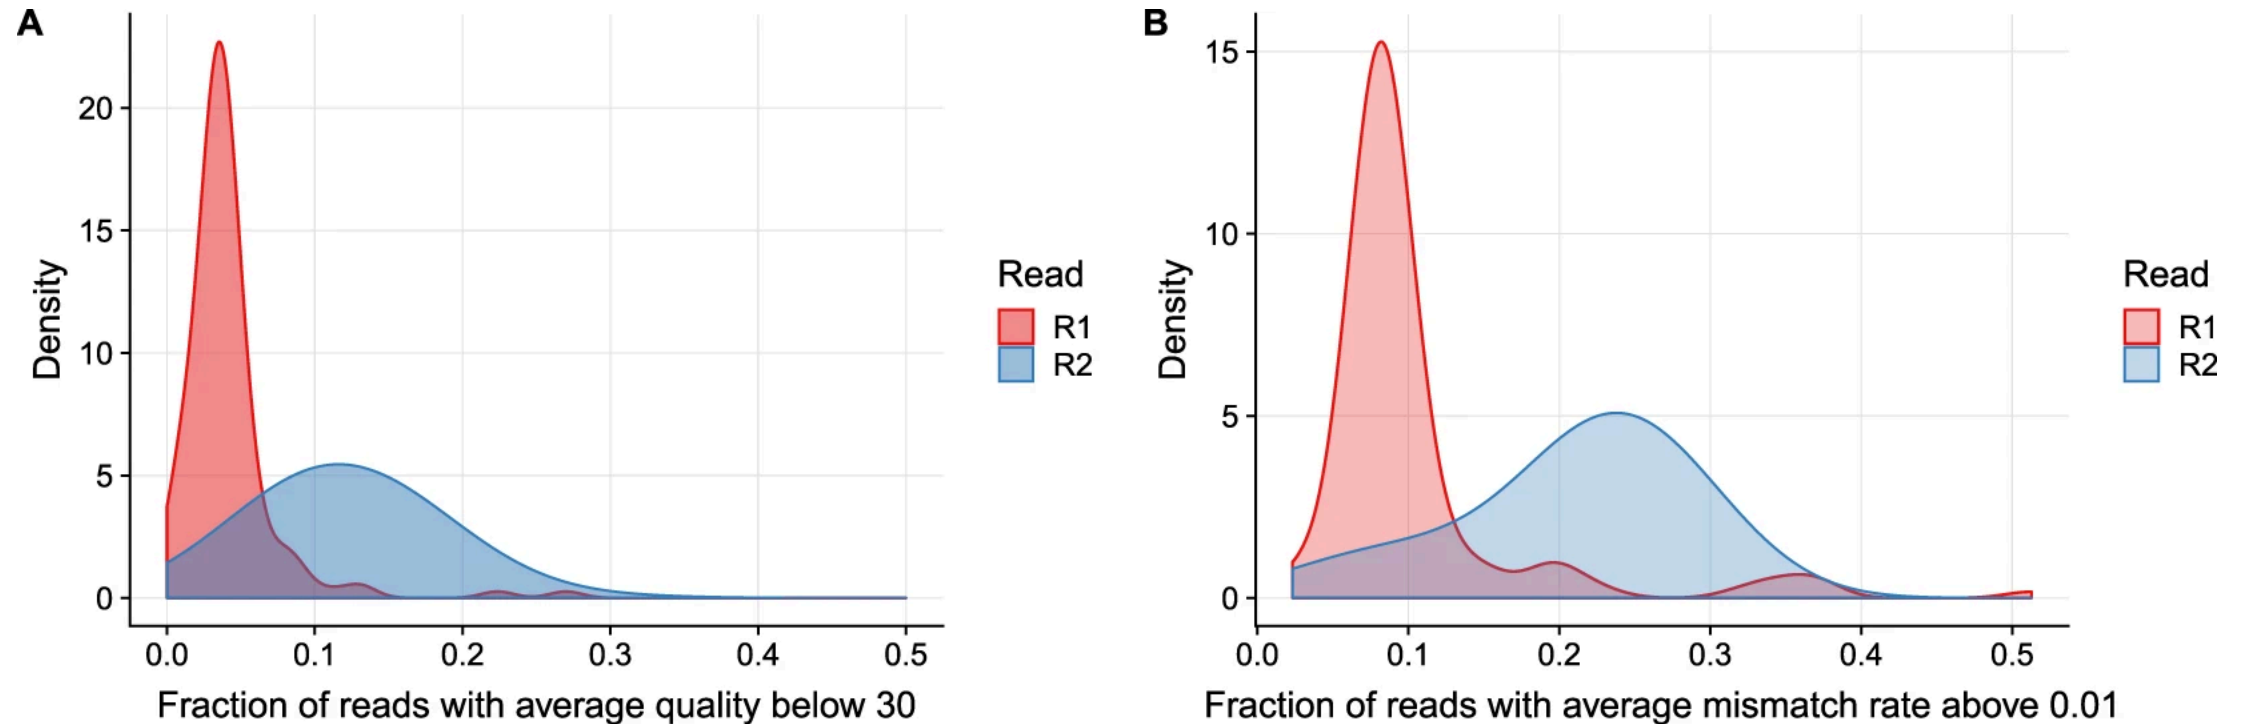
We want the SNPs to be distributed evenly across cycles

SNPs coming from towards end of the read are sign of false positive



# Long fragments have lower base quality

From: [Long fragments achieve lower base quality in Illumina paired-end sequencing](#)

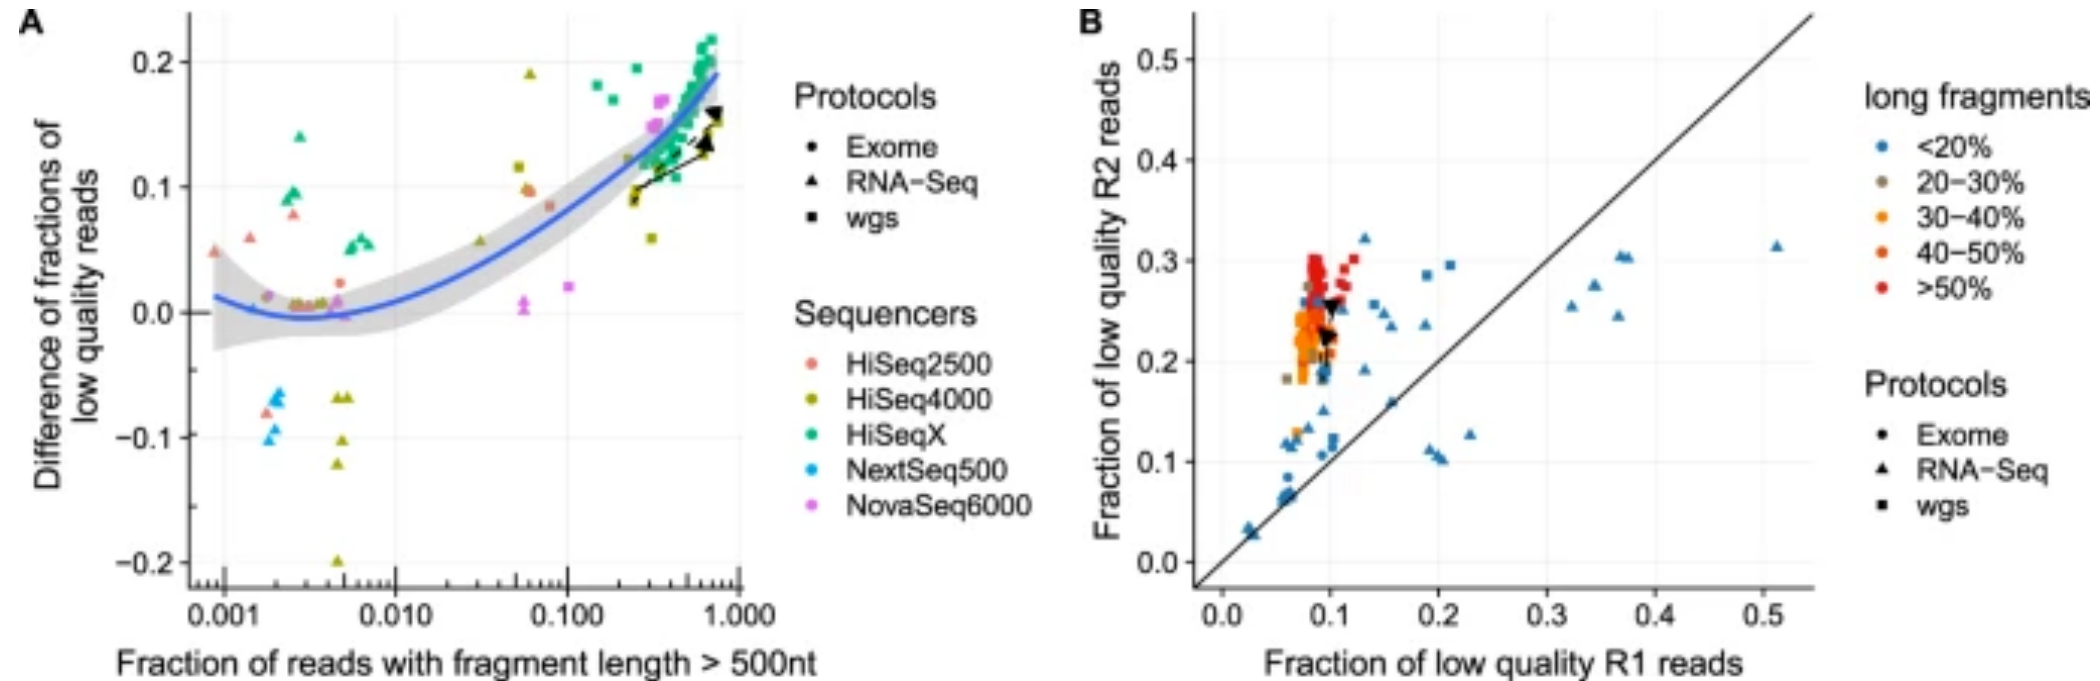


We plot the fraction of low quality reads in the 138 samples analyzed in our study. Across all samples the R2 reads harbor more low quality reads than the R1 reads. We plot two alternative definitions of ‘low quality’. Reads are called low quality if **(A)** the average Phred score is below 30, or **(B)** the average mismatch rate of the aligned bases is above 0.01. Both plots show that the R2 reads harbor more low quality reads and that the fraction of low quality reads is more variable across samples.



# Increase of R2 low quality reads as a function of the content of long fragments

From: [Long fragments achieve lower base quality in Illumina paired-end sequencing](#)



Increase of R2 low quality reads as a function of the content of long fragments. In **(A)** we plot for individual samples the difference in low quality read content among the R2 and the R1 reads versus the content of long fragments. The plot shows that the more long fragments a samples has the more prevalent are low quality reads among the R2 reads. In **(B)** we directly compare the fraction of low quality reads in R2 and R1 and color-code the content long fragments. Low quality reads are defined as reads having a mismatch rate above 0.01 in the bases after alignment. The plotted samples have been generated using various protocols on various sequencers in various labs. The dashed lines connect three samples each that have been processed identically except with an increasing targeted fragment length.




# Per base sequence content

- Sequence content across bases should not change with cycle


# Per base sequence content – RNAseq – typical Illumina library

- The primers used in the library are typically not removed




# Per base sequence content – targeted sequencing

- In targeted sequencing there is much less genes being sequenced so the base composition of reads is non-random



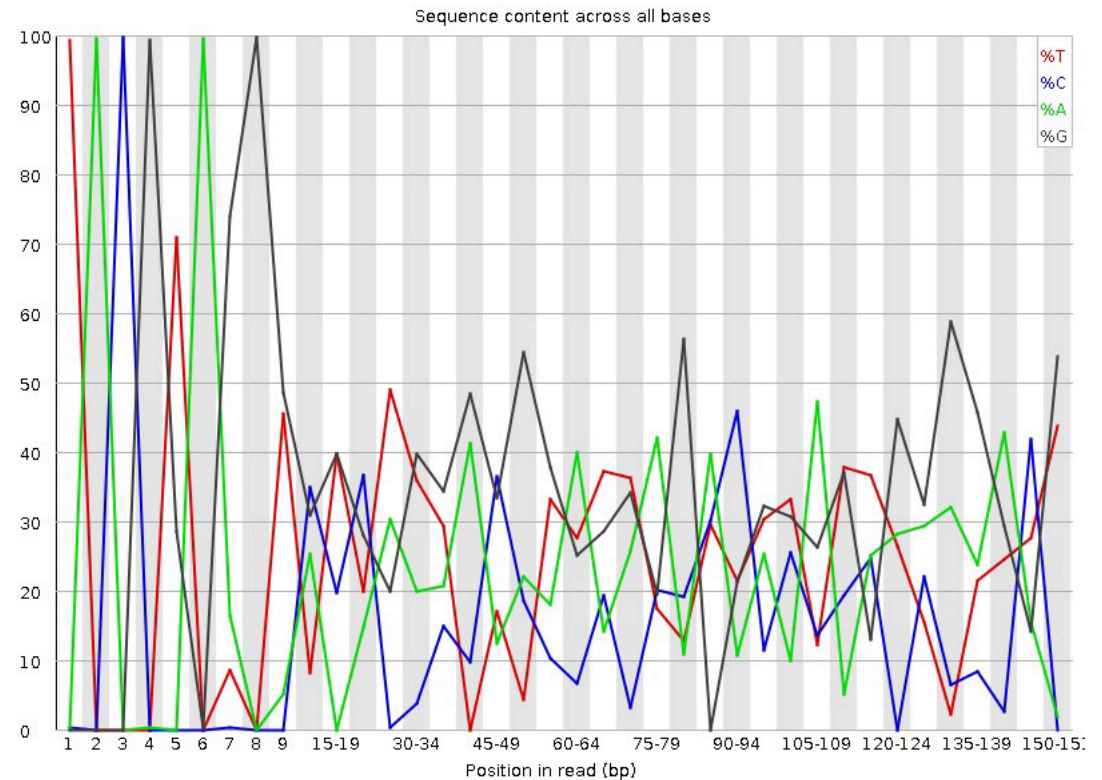
# Per base sequence content – a bad example?

- This suggests that a single sequence makes up a large part of the library – this can mean rRNA contamination in RNAseq



# Per base sequence content - microbiome

- ... however, it is expected if we sequence 16S rRNA of microbiome where one or few bacteria strains are dominating



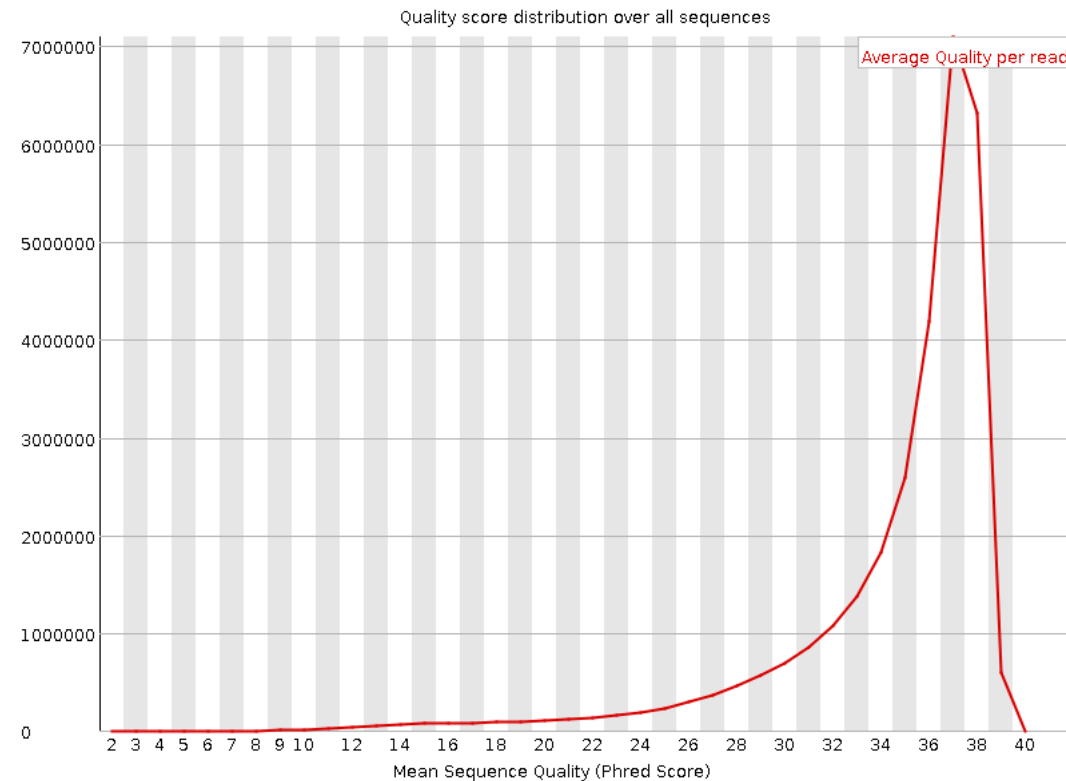


# Per sequence quality

- All – or at least majority of the sequences should have good average quality (average Phred score across all read bases)

# Per sequence quality - RNAseq


- majority of the sequences have good average quality






# Per sequence quality - microbiome

- Small peaks in lower average quality can suggest low quality ends on part of sequences – attention, if small read diversity (e.g. microbiome), this can be due to highly duplicated reads due to too deep sequencing



# Per sequence quality – targeted sequencing

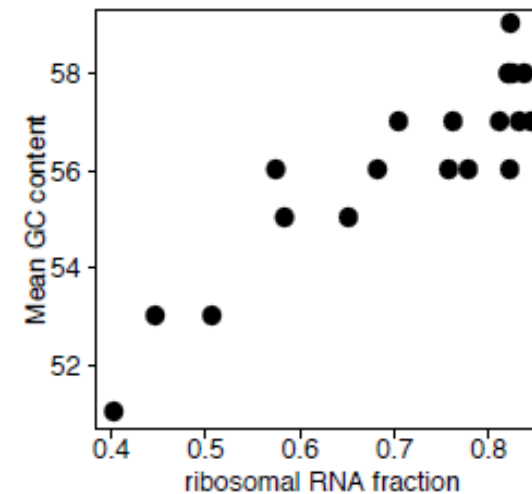
- Small peaks in lower average quality can suggest low quality ends on part of sequences – attention, if small read diversity (e.g. microbiome), this can be due to highly duplicated reads due to too deep sequencing



# Per sequence and per read GC content




- Mean GC content across reads should correspond to the overall GC content of the genome
- Even small shifts can indicate contamination with GC rich sequences (ribosomal RNA with high GC content for instance)




# Per sequence GC content - RNAseq

- A relatively good example of GC content




# Per sequence GC content – targeted sequencing

- This strange theoretical distribution is due to high amount of NNNNNN sequences in the reads




# Per sequence GC content – targeted sequencing after trimming

- The GC count per read is disturbed because of small number of genes sequenced!




# Per sequence GC content – microbiome

- The GC count per read is disturbed because of small diversity of sequences



# Per read GC content – good example


- The GC count per read is disturbed because of small diversity of sequences






# Per read GC content – typical RNAseq

- GC content different in first 8-10 bases, due to presence of primers




# Per read GC content – targeted sequencing

- GC content across different base positions due to high duplication level of reads and small diversity



# Per read GC content – microbiome

- GC content across different base positions due to high duplication level of reads and even smaller diversity.
- Zero GC in first two bases can be due to adapters.






# Per base N content

- In ideal case, there should be minimum of N calls in the reads
- “The HiSeq2000 produces very few Ns. It is very rare to see N content greater than 30%. When Ns are produced it is usually the result of some temporary instrument issue. For example a small bubble in the flow cell may cause focus problems at a certain cycle. Downstream processing of Ns depends on your analysis software and strategy.”
- Source:

[https://www.biotech.wisc.edu/services/dnaseq/sequencing/Illumina\\_old/Illumina\\_QC\\_FAQs](https://www.biotech.wisc.edu/services/dnaseq/sequencing/Illumina_old/Illumina_QC_FAQs)


# Per base N content – ideal case

• ...



# Per base N content – targeted sequencing

Enrichment for N calls at the beginning of the sequence






# Sequence duplication level and overrepresented sequences

- Indicates the library complexity and possible contamination
  - (the less duplicates, the more complex)
- Too many duplicated sequences means we sequenced “too much”.
- Overrepresented sequences may indicate:
  - Presence of adapters, presence of contamination (rRNA), PCR problems
- This holds, however, mainly for WGS, WES or RNAseq

# Sequence duplication level – good example (RNAseq)


- Most sequences occur only once





# Sequence duplication level – bad example (RNAseq)?

- Over-amplification! May come from highly expressed transcripts.



# Overrepresented sequences


## Overrepresented sequences

- Indicate remaining adapters, PCR duplicates, but also can be real sequences!
- Always judge based on type of data and check before filtering!

### Overrepresented sequences

Sequence	Count	Percentage	Possible Source
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNN	110195	6.337201578736401	No Hit
TGCACAGGAGCGGCTTCTGCAACTTCATGCATTTGAAGCCATTTCCAG	8442	0.4854907729723917	No Hit
ACCTTGGGAGGGTTTCAAGAAAGGCTGCATGTCCCAAGGCCATGCTAAC	7426	0.427061653647593	No Hit
CTCTATCTTCCCTAGTGTGGTAACCTCATTTCCCATAAAGATTCCAGAAC	7261	0.4175726726548846	No Hit
CCCCCTCAGCATCTTATCCGAGTGAAGGAAATTTGCGTGTGGAGTAT	7196	0.4138345892335147	No Hit
TGAGCAGGAGGGAAAAGTCTAATTACCATGACAAGAACATTGTATTAC	6902	0.39692695037377956	No Hit
CGTCTCCTGTTTTGTAGTCCAACCTGTGATGATTGATGCCAAGAAGTG	6668	0.3834698590568476	No Hit
GTCTTCATCTTATTGATGTTTTGATGGTCTTCTTATCCAACACGCCGAG	6521	0.3750160306269801	No Hit
TGCAAGCTCCTGGTGGCAGCTCTGAACGTATTTAAAAAAAATGAAATG	6359	0.36569957656141183	No Hit
TGCCCTGGCCCTGGCTTGTGGGGCTGCCAGCAGCTGCCATAAAGGAC	6352	0.36529701373141815	No Hit
CTGCCCCGAGGAGCACTAAGCGAGGTAAGCAAGCAGGACAAGAAGCGGT	6189	0.3513223326330657	No Hit
CCAGATGTTCTTCGCTAATAACCACGACCAGGAATTTGAGTCTGGGC	6070	0.3490794825802437	No Hit
CCTGTGTTATCTCCTAGTTGGCTCTGACTGTACCACCATCCACTACAAC	6065	0.3487919377016768	No Hit
GGCTCGGCACGCGCTACACACCTACTGCCGCCCTACCCCGGCTC	6058	0.3483893748716831	No Hit
TTCTCTGGAAACTCCCATTTGAGATCATATTATCTCTGAAATCAA	5910	0.3398780464661022	No Hit
TGCTCATGCCACAGAGACTTGACAACTGCAGAAATGGCAGCATTGG	5905	0.3395905015875353	No Hit
AAAGATGGAAGAGAAAGAGGATGGGTGGGAACTGTGCTCCCAT	5895	0.33901541183040146	No Hit
TCTCGAGGAGGAGTGACAGCAATGGCAGTTACTGTCAACAGGTGGACAT	5773	0.3319993167933685	No Hit
CTGGGTCTCCTCTTTCGGTCAAGGACTCTTTGCCAAGTTCACAGA	5672	0.3261909102463167	No Hit
ACATCCTGCTTACATCCTGGCAGGTACGGATCTAAACAGCGACTTTTT	5574	0.320555030626405	No Hit
CCTGGGACCCGATGCCTCTCTGCTGAGATCCCTCAGTTTTTCCCAG	5554	0.31940485111213734	No Hit
AGTGAGTGCAGTTGTTTACCATGATAACGACACACAAAATAGCCGTA	5511	0.3169319651564618	No Hit
AAAGATGGAAGTCCACCCCTTTCCTGGTTTTAAGTATGTATGGAATGTTA	5500	0.31629936642361456	No Hit
ACTGGAAGAAATGGATTCCAAGAGCAGTTCTCTCTTTAGTTGTGAAG	5425	0.31198619324511073	No Hit
GAGCTATGAGCTACGCCGCCCTCCCGATGTGGAGGTATGACTTCC	5424	0.31192868426939735	No Hit
AACCCACCAATTTTGGTAGCAGTGGAGAGCTACAGGACAACATGCCAGCA	5355	0.3079605649451738	No Hit
CCCATCCTCACCATCATCACACTGGAAGACTCCAGGTGAGGACCTTGCCACCTGCACACTGG	5317	0.3057752238680652	No Hit

# Sequence duplication level – targeted sequencing



Some reads are present more than 10 times. This is due to NNNN sequences and due to few genes sequenced (longer genes get more reads)

## Overrepresented sequences

Sequence	Count	Percentage	Possible Source
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN	110195	6.337201578736401	No Hit
TGCACACGAGGCGGCTTCTGCAACTTCATGCATTTGAAGCCATTTCCAG	8442	0.4854907729723917	No Hit
ACCTTGGGAGGTTTCAGAGAAGGCTGCATGTCACCAAGGCCCATGCTAAC	7426	0.427061653647593	No Hit
CTCTATCTTCCAGTGTGGTAACCTCATTTCCCATAAAGATTGAGAAC	7261	0.4175726726548846	No Hit
CCCCCTCAGCATCTTATCCGAGTGGAAAGAAAATTTGGGTGTGGAGTAT	7196	0.4138345892335147	No Hit
TGAGCAGGAGGGGAAAAGTCTAATACCATGACAAGAATTCATTATTAC	6902	0.39692695037377956	No Hit
CGTCTCCTGTTTGTAGTCCAACCTGTGATGATGATGCCAAAGAAGTG	6668	0.3834698508568476	No Hit
GTCTTCATCTTATTGATAGTTTTGATGGTCTTCTTCCAAACACGCCGAG	6521	0.3750168306269801	No Hit
TGCAAGTCTCTGGTGGCAGCTCTGAACGGTATTTAAAACAAAATGAAATG	6359	0.36569957656141183	No Hit
TGCCCTGGCCCTGGCTTGTGGGGTCCCCAGCAGTCCCATAAAGGAC	6352	0.36529701373141815	No Hit
CTGCCCCAGGGAGCCTAAGCGAGTAAGCAAGCAGGACAAGAAAGCGGT	6109	0.3513223326330657	No Hit
CCAGATGTTCTTCGTAATAACCACGCCAGGAATTTGTGAGTCTGGGC	6070	0.3490794825802437	No Hit
CCTGTGTATCTCTAGTTGGCTCTGACTGTACCACCATCCACTACAAC	6065	0.3487919377016768	No Hit
GGCTGGGCACGCGCTACACACCTACTGCCCGCCCTACCCGGCTC	6058	0.3483893748716831	No Hit
TTCTCTGGAACTCCATTTGAGATCATATTCATATTCTCTGAAATCAA	5910	0.3398780464661022	No Hit
TGCTCATGCCACAGAGACTGCACAACATGCAGAATGGCAGCACATTGG	5905	0.3395905015875353	No Hit
AAAGATGGAAGAGAGAAAGGCATGGGTGGGAACTGTGCCTCCATT	5895	0.33901541183040146	No Hit
TCTGAGGAGGCAGTGACAGCAATGGCAGTTACTGTCAACAGGTGGACAT	5773	0.3319993167933685	No Hit
CTGGGCTCTCTCTTTGTTCAAGGACTCTTTGCCAAGTTCACAGA	5672	0.3261909102463167	No Hit
ACATCTGCTTACATCTGGCAGGTAGCATTAACAGCGACTTTTTT	5574	0.320555030626405	No Hit
CCTGGGACCAGGATGCTCTCTCTGCTGAGATCCCTCCAGTTTTTCCAG	5554	0.3194048511213734	No Hit
AGTGAGTGCAGTTGTTCATGATAACGACACACAATAAGCCGTA	5511	0.3169319651564618	No Hit
AAAGATGGAAGTCCACCTTTGCTGGTTTTAAGTATGATGGAATGTTA	5500	0.31629936642361456	No Hit
ACTGGAAGAAATGGATCCAAAGAGCAGTTCTCTCTTTAGTTGGAAG	5425	0.31198619324511073	No Hit
GAGCTATGAGCTACGGCCGCCCTCCCGATGTTGGAGGTATGACCTCC	5424	0.31192868426939735	No Hit
AACCCACAAATTTTGGTAGCAGTGGAGAGCTACAGGACCACTGCCAGCA	5355	0.3079605649451738	No Hit
CCCATCTCACCATCATCACTGGAAGACTCCAGGTCCAGGACCACTTGCCACCCTGCACACTGG	5317	0.3057752238680652	No Hit

# Quality control exercise

- We continue in our exercise from `1_Preprocessing.sh`

## Step 1: Read quality control and **data filtering**

- Based on the quality measures, we decide to remove low quality bases and reads

- **Trimming** – removes low quality or unwanted bases from reads, thus shortening them. Is applied to increase the number of mappable reads.
- **Filtering** – removes whole reads that do not meet quality standards (e.g. too short etc)



# Trimming reads

- Read trimming is applied to increase the number of mappable reads by:
    - Removing low quality bases at the end of the reads that are likely to contain sequencing errors
    - Removing adapter sequences
-

# Removing adapters

- Important mainly for very short read sequences of interest (when the input DNA fragment is less than the read length)  
e.g. for miRNA with 22nt length the adapter gets sequenced more often than for RNA sequences, which are much longer

What is the sequence of adapters?

Best option: ask which kit was used for preparing libraries

Programs: `cutadapt`, `trimmomatic`

TruSeq Universal Adapter: 5

AATGATACGGCGACCACCGAGATCTACACTCTTTCCCTACACGACGCTCTTCCGATCT 3

TruSeq Indexed Adapter 5

GATCGGAAGAGCACACGTCTGAACTCCAGTCAC NNNNNN ATCTCGTATGCCGTCTTCTGCTTG 3

Here “N” is any nucleotide, and the 6 of them together are a unique sequence which can readily be identified as unique to 1 library.

---



# Filtering reads

- We can remove whole reads based on:
  - quality of its base calls
  - its length (too short reads)
  - level of duplication
  - ...
-





# Trimming and filtering - exercise

## Trimming and filtering - exercise

- We continue in our exercise from `1_Preprocessing.sh`
  - We will use `grep` command to find adapter sequences and `cutadapt` to remove them
  - We will trim low quality bases
  - Independent work: find specific QC problems in your project data and suggest solutions (what to trim, filter, etc)
-

# Recommended literature

- Fuller et al. 2009: The challenges of sequencing by synthesis  
[http://arep.med.harvard.edu/pdf/Fuller\\_09.pdf](http://arep.med.harvard.edu/pdf/Fuller_09.pdf)
- <https://sequencing.qcfail.com/>