

Introduction to RNAseq


Version 1

Contents

Introduction to RNAseq	2
1 Genomics	4
1.1 Sequence	4
1.2 Read sequence	5
1.3 Contig sequence	6
1.4 Summary	6
2 Bioinformatics file formats	8
2.1 Introduction	8
2.2 What is a file format?	9
2.3 Annotation files	9
2.4 Glossary	10
3 Reference sequence	11
3.1 FASTA	11
3.2 Exercise	12
3.3 Rule of thumb	13
4 Gene annotation	14
5 Raw data	15
5.1 Introduction	15
5.2 FASTQ	15
5.3 Unix tooling	18
6 RNAseq bioinformatics workflow	19
6.1 Adapter trimming	19
6.2 Alignment	19
6.3 Read counting	21
6.4 RNAseq checklist	24
6.5 Summary	25
7 QC	27
7.1 Number of mapped reads	27
7.2 STAR aligner	28
7.3 Number of assigned reads	29
7.4 FastQC	30
7.5 Intragenic vs Intergenic region	30
7.6 Reads mapping bias	30
7.7 Insert size distribution	30
7.8 Where to look for more help	30
8 All about RNAseq	31
8.1 General	31
8.2 Technology	31
8.3 Differential expression (DE)	31

9 Lesson plan	33
10 Meeting notes	34

Introduction to RNAseq



These are course notes for the “Introduction to RNAseq” course given by the Monash Bioinformatics Platform¹ for the Monash Data Fluency² initiative. Our teaching style is based on the style of The Carpentries³.

- PDF version for printing⁴

Source code

- GitHub page⁵

Authors and copyright

This course is developed for the Monash Bioinformatics Platform by Stuart Archer, Nick Wong and Kirill Tsyganov.

¹<https://www.monash.edu/researchinfrastructure/bioinformatics>

²<https://monashdatafluency.github.io/>

³<https://carpentries.org/>

⁴<https://monashdatafluency.github.io/rnaseq-intro/rnaseq-intro.pdf>

⁵<https://github.com/MonashDataFluency/rnaseq-intro>



This work is licensed under a CC BY-4: Creative Commons Attribution 4.0 International License⁶. The attribution is “Monash Bioinformatics Platform” if copying or modifying these notes.

⁶<http://creativecommons.org/licenses/by/4.0/>

Chapter 1

Genomics

Genome is a collection of all genetic material in the organism. To be more accurate each cell in our body contains a copy of entire genome. The study of genomics is all things to do with genome including gene expression, methylation profiles, transcriptional factor binding, chromatic arrangements and three dimensional structure of chromosomes and many more things.

In this course we are going to focus on one particular technique of genomics - RNA sequencing or RNAseq for short. This is a study of primarily but not limited to gene expression. Other common application of RNAseq are study of splicing and isoforms.

In genomics everything rotates around sequence information, let's look into biological sequence next.

1.1 Sequence

A sequence is a generic name describing the order of biological letters in molecules like DNA, RNA or ptiptide and protein. In the case of DNA and RNA letters mean nucleotide bases, **A**denine, **T**hymine, **C**ytosine **G**uanine or ATCG for short. Also know that mRNA contains **U**racil instead of Thymine. This would be one way to represent all five possible bases A[T|U]CG In the case of ptiptide/protein molecules each letter means amino acid type, I won't list discuss proteins any further in this book since it isn't a focus on this book.

Note that there is a greater nomenclature exists to cover all possible variances

##	Letter_code	Base_name
## 1	A	Adenine
## 2	C	Cytosine
## 3	G	Guanine
## 4	T U	Thymine Uracil
## 5	R	A G
## 6	Y	C T
## 7	S	G C
## 8	W	A T
## 9	K	G T
## 10	M	A C
## 11	B	C G T
## 12	D	A G T
## 13	H	A C T
## 14	V	A C G
## 15	N	any base

The reason we have sequence information of any kind is because someone somewhere made an afford to sequence some part of the molecule. In the case of whole genome sequencing an afford was made to sequence the entire genome, for human the entire genome is 3 billion bases long. However this is a sum of

all bases across all chromosomes, but chromosomes are not all equal, below is a sequence length of each chromosome.

```
## Chrom Length
## 1 1 248956422
## 2 2 242193529
## 3 3 198295559
## 4 4 190214555
## 5 5 181538259
## 6 6 170805979
## 7 7 159345973
## 8 8 145138636
## 9 9 138394717
## 10 10 133797422
## 11 11 135086622
## 12 12 133275309
## 13 13 114364328
## 14 14 107043718
## 15 15 101991189
## 16 16 90338345
## 17 17 83257441
## 18 18 80373285
## 19 19 58617616
## 20 20 64444167
## 21 21 46709983
## 22 22 50818468
## 23 X 156040895
## 24 Y 57227415
## 25 MT 16569
```

While sequence is a generic name and often times it needs some context around it, for example DNA sequence, RNA sequence or chromosome 3 sequence. There are other names like contig and a read, which are both sequences too, they mean more specific thing. In general both mean shorter sequence, where contig almost always should be longer than a read.

Let's talk about a read sequence next.

1.2 Read sequence

A read is product of a machine read out that takes some biological material i.e DNA or RNA sample and produces number of reads that resembles what you've put into the machine. There have been a few iteration of sequencing technologies, each iteration brings about higher yield (more reads) and longer reads. Current players in the genomics space are:

Short read technologies

- Illumina¹
 - MiSeq
 - HiSeq
 - NextSeq
 - NovaSeq
- BGI platforms²
 - BGISEQ-50
 - BGISEQ-500

¹<https://www.illumina.com/>

²<https://www.bgi.com/global/resources/sequencing-platforms/>

- Ion Torrent³

Medium to long read technologies

- PacBio⁴
 - Sequel System
- Oxford Nanopore⁵
 - MinION
 - GridION
 - PromethION

All of the above technologies will one way or another result in digital reads. As name implies short reads are typically 50-300 bases, whereas long reads can range from 1k to 10k or more. Usually the quantity and the length of the read are inversely proportional, however the throughput of a short reads technology at the moment is vastly greater than the long read and it doesn't suffer from the coverage issue. In terms of number of reads we are talking about millions to billions for short reads technology and in the range of 1 to 4 millions for long reads. This is a good plot showing nanopore output, which illustrates quantity vs length nicely.

Getting the reads is just a first step. You have to do something with your reads. The two paths one can take are really assemble your reads into contigs and further into full genome or map your reads to an existing reference genome. In this course we'll be taking the later path - mapping our RNAseq reads to an existing reference genome. But let me just talk a little more about what a contig sequence is next.

1.3 Contig sequence

Contig has specific meaning in genome assembly. It is a larger fragment assembled from short (100-300 bp) reads that hasn't yet been merged into chromosome. However often times this term can be loosely used to mean any long stretch of sequence. I often (incorrectly) refer to chromosomes as different contigs in the FASTA file. With long read technologies advancing pretty rapidly it breaks conventional definitions of a read and contig, where a single read can be longer than a contig. However assembling two long reads would still result in a contig sequence. So once again it is important to give more context when using this term.

1.4 Summary

- **Genomics** is all things genome that rotates around genomic sequence
- **Sequence** is an ordered collection of nucleotide bases
- **Read** is an output from a sequencing machine that covered biological sample in digital sequence
- **RNAseq** focuses on gene expression and splicing isoforms

³<https://www.thermofisher.com/au/en/home/brands/ion-torrent.html>

⁴<https://www.pacb.com/>

⁵<https://nanoporetech.com/>

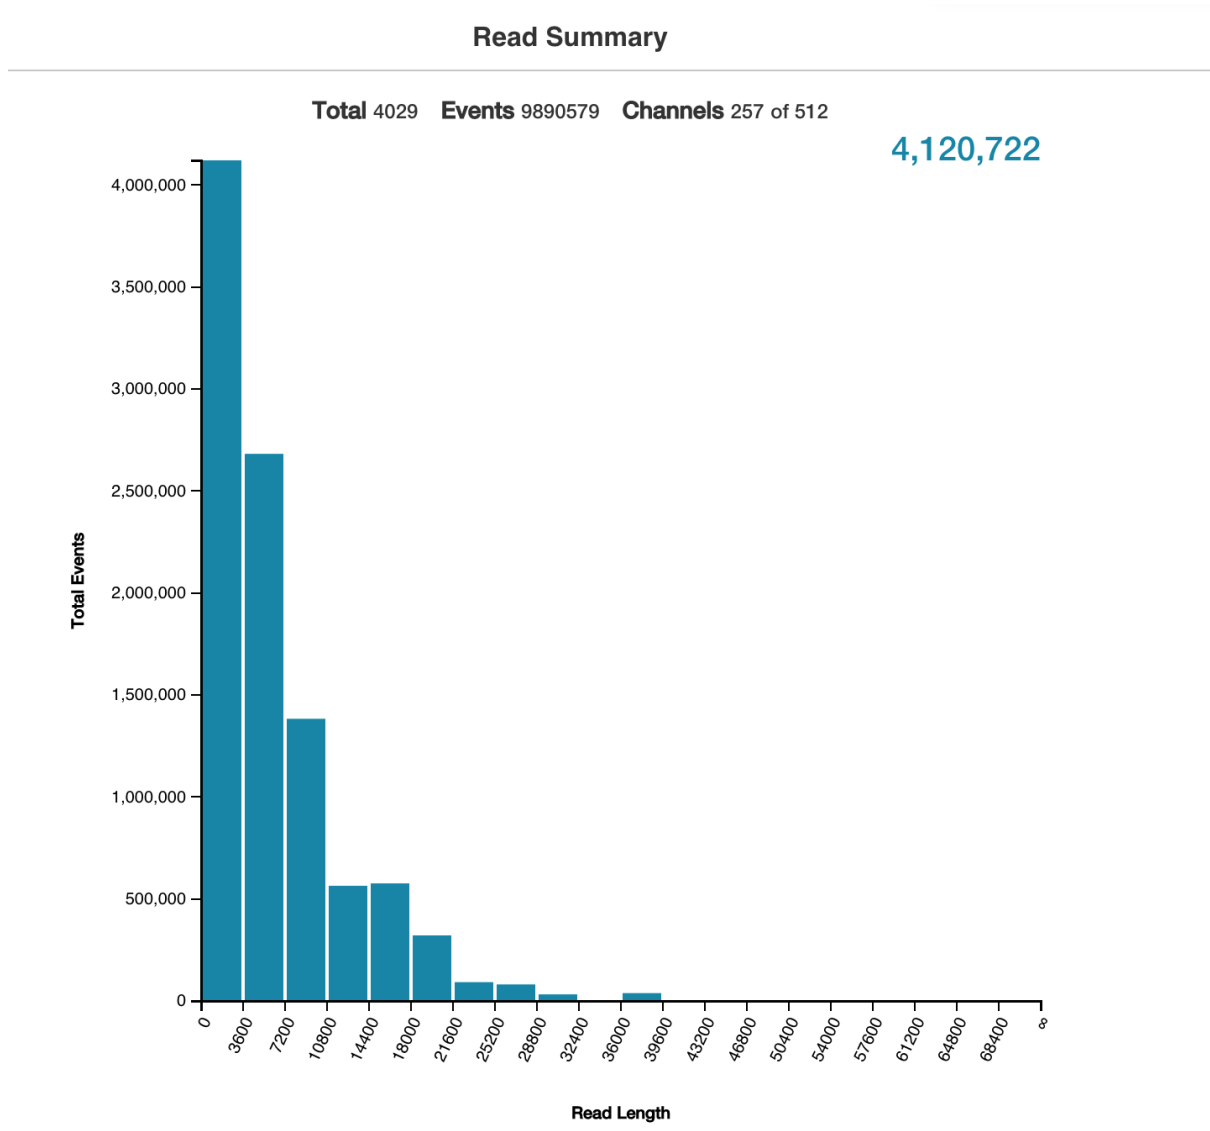


Figure 1.1: <https://bioinformatics.stackexchange.com/questions/3623/how-many-reads-has-my-sequencing-run-produced-on-minion>

Chapter 2

Bioinformatics file formats

Before diving into specific file formats. I would like to discuss what could file actually mean and hold in general, after all bioinformatics files aren't that different from any other files. In fact fair number of bioinformatics files are just a variant of TSV file format, discussed shortly. Bioinformatics world largely rotates around Unix/Linux ecosystem, and remember that macOS is Unix derived operating system too. That means that a lot of the tooling and compute happens on linux computers often referred to as servers. In unix world file name is simply a string of letters (text) that points to a location of that file content that is stored elsewhere. Therefore the only thing that matters for file names is they are unique. Often people get caught out by file's extensions, they do not matter. File extensions are simply an indicator for a user what to expect when they are going to open that file. To more accurate other programs often look at file extensions too trying to guess if this is the file format they can handle. In general it is a very good idea to have correct file extension. I just want you to know that just because you've renamed a file from `ref.fasta` to `ref.fa` nothing had happened to its content. Bioinformatics is an evolving field and new technologies will bring about new file formats.

2.1 Introduction

Broadly speaking any computer file is either a plain text or binary format. Under the hood everything is binary of course, but a plain text file can be looked at and modified using standard unix/linux tool sets as well simply opened with any text editor for more custom modification. If the file isn't plain text you have to have specific tool to work with that file type. For example pdf files aren't plain text and this is why you need pdf viewer (e.g adobe reader) to work with pdfs. Similarly in bioinformatics BAM files aren't plain text and do require `samtools` to work with it. As I've hinted earlier large number of bioinformatics files are plain text, however a lot of the times plain text files are zipped with `gzip` tool adding another suffix to the file name and thereby changing the file into binary for that now will require `gzip` tool to work with. We zip file to reduce their size. Zipping works particularly well on files that have repetitive patterns e.g repeats of ATCG sequence.

Note the difference in size is essentially three times !

```
-rw-r--r-- 1 kirill kirill 3.0G Apr 10 10:11 Homo_sapiens.GRCh38.dna_sm.removed_contigs.fa
-rw-r--r-- 1 kirill kirill 895M Apr 10 10:10 Homo_sapiens.GRCh38.dna_sm.removed_contigs.fa.gz
```

It is a natural behaviour wanting to unzip file before looking inside. It is okay to do so, but be aware that this isn't the best practice, and you don't want to be storing your files in plain text for a long time, this will eat up your disk space very quickly.

example of large BAM file and plain text SAM

Remember that best practice is always keep your files zipped

Let's talk next about what makes different file formats different

2.2 What is a file format?

The file format at its simplest defines structure within the file. By far the most common structure of the plain text file is tabular. The spreadsheet that everyone is so familiar with is simply a plain text file with rows and columns, where each cell separated by a “special” character called **tabular** or **tab** for short, which makes tab separated variables file or tsv for short with a sister file of comma separated variables or csv for short where instead of **tab** character comma is used.

Tab Separated Variables (tsv)

Gene.ID	Chrom	Gene.Name	Biotype	K03_S3	K04_S4	WT1_S1	WT2_S2
ENSMUSG00000051951	1	Xkr4	protein_coding	0	0	0	0
ENSMUSG00000025900	1	Rp1	protein_coding	0	0	0	0
ENSMUSG000000109048	1	Rp1	protein_coding	0	0	0	0
ENSMUSG00000025902	1	Sox17	protein_coding	0	0	0	0
ENSMUSG00000033845	1	Mrpl15	protein_coding	934	1317	888	939
ENSMUSG00000025903	1	Lypla1	protein_coding	705	848	647	747
ENSMUSG000000104217	1	Gm37988	protein_coding	0	0	0	0
ENSMUSG00000033813	1	Tcea1	protein_coding	354	436	368	442
ENSMUSG00000002459	1	Rgs20	protein_coding	0	0	0	0

Comma Separated Variables (csv)

Gene.ID	Chrom	Gene.Name	Biotype	K03_S3	K04_S4	WT1_S1	WT2_S2
ENSMUSG00000051951	1	Xkr4	protein_coding	0	0	0	0
ENSMUSG00000025900	1	Rp1	protein_coding	0	0	0	0
ENSMUSG000000109048	1	Rp1	protein_coding	0	0	0	0
ENSMUSG00000025902	1	Sox17	protein_coding	0	0	0	0
ENSMUSG00000033845	1	Mrpl15	protein_coding	934	1317	888	939
ENSMUSG00000025903	1	Lypla1	protein_coding	705	848	647	747
ENSMUSG000000104217	1	Gm37988	protein_coding	0	0	0	0
ENSMUSG00000033813	1	Tcea1	protein_coding	354	436	368	442
ENSMUSG00000002459	1	Rgs20	protein_coding	0	0	0	0

Common bioinformatics files such as gff/gtf, bed and vcf are all variants of tsv in some respect except with defined column names and hence value types that can fit in those columns. I haven’t explained what each of those files is used for. Let’s focus on gff/gtf format only and leave the other two for later discussion.

2.3 Annotation files

The essence of an annotation file, guessable from the name, is to provide information of coordinates of features in the genome. A feature could mean several things including a gene, transcript, exon, 3’UTR or any other region of the genome that is “worth” documenting. Basically this is your google map without pictures. If you think about what information you would need in order to identify a gene for example well known tumor suppressor gene TP53.

- bear minimum we need to know which chromosome it is on
- where about it is on chromosome 17
- remember that humans are diploid - which strand of chromosome 17 it is on

In a simple tabular format this would look like this

```
tp53    chr17    7661779 7687550 +
```

Turns out that this is actually an annotation file format called Simplified Annotation Format (SAF)¹. It isn’t that widely used, because it has some limitations, which is outside of the course to discuss and there are at least three other annotation formats that exist, which we are not going to discuss any further. Side tracking to our earlier discussion about file extensions, **saf** is a file format given that it has defined structure, but what should the extension be? Anything you like **.txt** because it is a text file, **.tsv** well it

¹<http://bioinf.wehi.edu.au/featureCounts/>

is tab separated. I usually do `.saf`, because that tells me straight away exactly what the content of the file is. This is what file extension supposed to do.

The file format that we are going to focus on is General Feature Format (gff)

- Plain text file format
- The format was proposed as a means to transfer feature information
- We do not intend GFF format to be used for complete data management of the analysis and annotation of genomic sequence
- Originally wanted file format that will be easily parsable by Unix tools such `grep`, `sort` and `awk`

2.4 Glossary

- interface: extract or insert relevant information

Chapter 3

Reference sequence

In this course will be focusing on DNA/RNA molecules and hence our sequences will be nucleotides sequences. In the case of *Mus_musculus* (mouse) genome, we have 19 autosomal chromosomes and two sex chromosomes X and Y plus a mitochondrial chromosome (genome). This makes up 22 separate molecules and therefore 22 separate sequences. Typically that sequence is stored in **FASTA** file format

- FASTQ for raw short reads
- FAST5 for raw long reads
- FASTA for any other kind of sequence

The other possible bioinformatics files that we weren't focus on are, genbank and gff. Both of those files can hold annotation and sequence information in one file. It often works fine with small genomes, but large eukaryotic genomes are better split into two files reference sequence and annotation information.

3.1 FASTA

There are a few different fasta files, but all of them will hold some sort of sequence. Sometime you'll see `.fna` to indicate nuclear sequence or `.faa` to indicate amino acid sequence, but more generally it'll be `.fa` or a longer version `.fasta`. Like any biological plain text file. This most certainly will be zipped so expect to see `.fa.gz` or `.fasta.gz`.

For the purpose of RNAseq analysis know that FASTA file holds your reference sequence.

FASTA file typically can hold any type of sequence data. We are not going to talk about amino acid sequence in this course and so for your purpose we can think of FASTA as just holding nucleotides sequence. That sequence can be anything, primer or adapter sequence, individual transcripts (genes) sequence OR entire chromosome sequence. In the case when FASTA holds entire genome, that is all chromosomes for a given species, we tend to refer to that FASTA file as a reference genome file. If I'm asking biologist of a **reference genome**, I typically imply that I need a FASTA file with genomic sequence.

Be aware that there isn't such thing as "the genome" for a particular species. This is because any genome is a

There are a few different fasta files, but all of them will hold some sort of sequence. Sometime you'll see `.fna` to indicate nuclear sequence or `.faa` to indicate amino acid sequence, but more generally it'll be `.fa` or a longer version `.fasta`. Like any biological plain text file. This most certainly will be zipped so expect to see `.fa.gz` or `.fasta.gz`.

Note that our purpose we can think of FASTA file holding our reference genome sequence of interest e.g mouse genome.

This is previously assembled genome that was made puplically available. There are several different vendors that provide reference genomes. In general either one of these three will is a good place to start looking for you reference genome.

- Ensembl¹
- RefSeq²
- UCSC genomes³

Those will have majority of species, however it is common to have a stand along community that host reference genomes from a different place. A good example is yeast genome database⁴. If you are working with yeast this is the place to look for reference genomes.

It is important to note that reference genome isn't a defined sequence. It is simply the best we have so far and it is a continuous effort to improve quality of the genomes. For widely used model organisms like mouse and yeast their genomes have been sequenced and assembled many times over the years and we have very good quality genomes. Species like axolotl or spiny mouse are less commonly used as model organism in comparison to mouse genomes and quality of those genomes is poorer. There are however other factors in place that influence genome quality e.g genome complexity.

3.2 Exercise

- go to Ensembl website⁵
- from the drop down on the left hand site select *Saccharomyces_cerevisiae*, which is a yeast species.
- click on download FASTA file. This will redirect you to ftp site
- find file **Saccharomyces_cerevisiae.R64-1-1.dna_sm.toplevel.fa.gz** and download it by clicking
- once downloaded double click on the file to open it and view the content.

depending on the operating system that you have you might have to unzip fasta file first before opening. It is best not to open with MS Office suite. Use text editor instead. On windows it'll be **notepad** on macOS it'll **textmate**.

At the very top you should see this

```
>I dna_sm:chromosome chromosome:R64-1-1:I:1:230218:1 REF
ccacaccacacccacacccacacaccacaccacacaccacacccacacacaca
caTCCTAACACTACCCTAACACAGCCCTAATCTAACCCCTGGCCAACCTGTCTCTCAACTT
ACCCTCCATTACCCTGCCTCCACTCGTTACCCTGTCCCATTCAACCATAACCACTCCGAAC
```

Note the very first line.

```
>I dna_sm:chromosome chromosome:R64-1-1:I:1:230218:1 REF
```

This is a FASTA header that indicates which sequence is follows. In this case it is I i.e sequence from chromosome one of yeast. If there are multiple chromosomes (usually the case for eukaryotic species) then you keep reading the sequence until you hit either the end of the file OR the next ">" sign

NOTE the actual sequnce have been truncated

```
>I dna_sm:chromosome chromosome:R64-1-1:I:1:230218:1 REF
ccacaccacacccacacccacacaccacaccacacaccacacccacacacaca
caTCCTAACACTACCCTAACACAGCCCTAATCTAACCCCTGGCCAACCTGTCTCTCAACTT
ACCCTCCATTACCCTGCCTCCACTCGTTACCCTGTCCCATTCAACCATAACCACTCCGAAC
>II dna_sm:chromosome chromosome:R64-1-1:II:1:813184:1 REF
CCCACACACCACCCACACCACCCACACACCACACACCACACCACCCACACCCACA
CCACACCACACCACACCACACCACCCACACCCACACCACACCACACCACACCACACC
ACACCACACACACCACACCACACCACACCACACCACACCACACCACACCACACCACACC
CACACCACACCACACCACACCACACCACACCACACCACACCACACCACACCACACCACACC
```

The format of the FASTA header is rather loosely defined, but these are key points:

- must start with ">" sign

¹<https://asia.ensembl.org/index.html>

²<https://www.ncbi.nlm.nih.gov/refseq/>

³<http://hgdownload.soe.ucsc.edu/downloads.html>

⁴<https://www.yeastgenome.org/>

⁵<https://asia.ensembl.org/index.html>

- the contig/chromosome name follows straight after “>” sign. Can be any string
- anything after a space is a comment/description

Below are all valid FASTA headers

```
>I
>chrI
>NC_000067.6
>I_named_it
```

Each header in a single FASTA file has to be unique in order to identify each sequence.

Are these two FASTA headers the same?

```
>I
>I named it
```

Are these two FASTA headers unique?

```
>I
>I_named_it
```

While these appear to be very small details, the reason it is important is because you want to know which chromosome (molecule) your reads came from.

3.3 Rule of thumb

Unless you are working with some exotic species and/or have a strong preference for genome vendor you can generally leave this part up to your bioinformatician to sort out. It is important to know that there are different vendors that have different versions (builds) of genomes.

. Often times bacterial RNAseq data come with some custom reference and annotation files. If that happens bioinformatician will need to convert everything to FASTA and GTF/GFF files.

Chapter 4

Gene annotation

As discussed in [reference genome] module. The reference genome is just a string of characters.

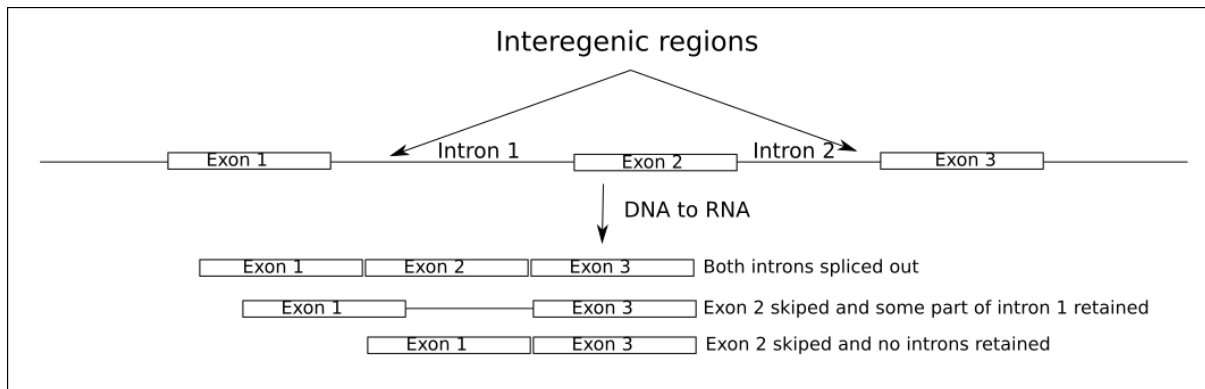


Figure 4.1:

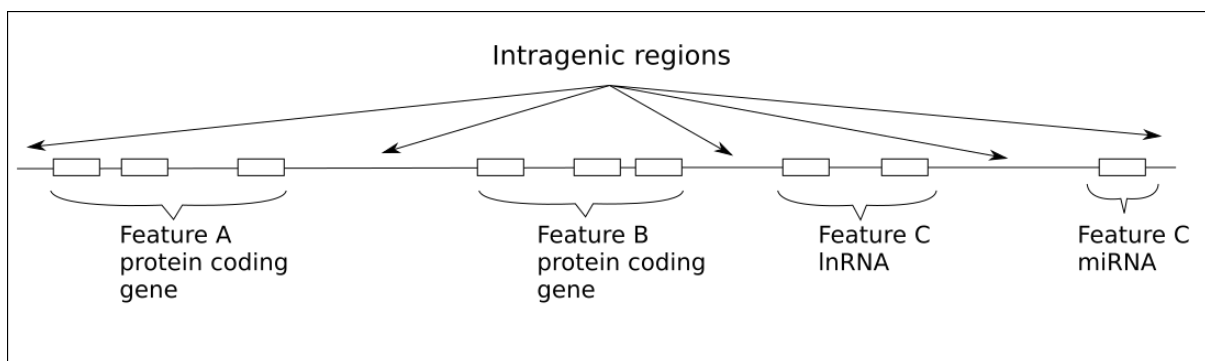


Figure 4.2:

Chapter 5

Raw data

5.1 Introduction

The raw data comes from a sequencing instrument, Illumina instruments appears to dominate short-read sequencing described earlier.

5.2 FASTQ

This is a very important file(s) you don't mess directly with FASTQ files. Best not to every unzip those either. If think that you need to do some thing fastq files in terms of editing or some other out of ordinary manipulations. Make a copy and work on the copy. If you modify, corrupt or loose your FASTQ files, days, month or year of wet-lab work go down a toilet, let a lone large sum of money spend on RNA extraction, library preparation and sequencing.

Your reads come in the form of FASTQ files. By far the most dominant way of organising your FASTQ files one per sample. For example if your experiment is a straightforward TP53 knock-out in mice. You'll need bear minimum 4 mice, best to get 6, this is 3 control and 3 tp53 knock-outs. This is way you have three replicates to account for mouse to mouse variability and have a better chance of estimating actual difference due to knock-out.

You should receive these files from sequencing facility:

```
cntr_rep1_R1.fastq.gz
cntr_rep2_R1.fastq.gz
cntr_rep3_R1.fastq.gz
```

```
tp53_ko_rep1_R1.fastq.gz
tp53_ko_rep2_R1.fastq.gz
tp53_ko_rep3_R1.fastq.gz
```

The naming of the files is completely up to you. Sequencing facility doesn't care about the names, as long as you give them 6 different microtubes with some unique naming on the lid, but remember that giving clear names simplifies the matter. Bioinformatician might be able to infer from FASTQ file names what the experiment was.

Question1: How many FASTQ files do you expect in a single-end experiment with 6 samples?

Remember that if you decided to sequence from both ends of the reads i.e paired-end sequencing, you'll give FASTQ file for each sample for each end. You should expect double the number of samples. It is important to have exactly double the number of files, because downstream analysis will rely on those files and will break if not supplied.

```
cntr_rep1_R1.fastq.gz
cntr_rep1_R2.fastq.gz
```

```
cntr_rep2_R1.fastq.gz
cntr_rep2_R2.fastq.gz
```

```
cntr_rep3_R1.fastq.gz
cntr_rep3_R2.fastq.gz
```

```
tp53_ko_rep1_R1.fastq.gz
tp53_ko_rep1_R2.fastq.gz
```

```
tp53_ko_rep2_R1.fastq.gz
tp53_ko_rep2_R2.fastq.gz
```

```
tp53_ko_rep3_R1.fastq.gz
tp53_ko_rep3_R2.fastq.gz
```

Question2: How many FASTQ files do you expect in a paired-end experiment with 6 samples?

One other thing that you should be aware is that sometimes to get extra deep sequencing your individual samples can be split across different lanes on the sequencing instrument (mainly talking about Illumina instruments here). This might seem confusing but you will need to multiple your samples by number of lanes. If you have single-end data of 6 sample that were split across 2 lanes then you are going to end up with 12 FASTQ files.

```
cntr_rep1_R1_L001.fastq.gz
cntr_rep1_R1_L002.fastq.gz
```

```
cntr_rep2_R1_L001.fastq.gz
cntr_rep2_R1_L002.fastq.gz
```

```
cntr_rep3_R1_L001.fastq.gz
cntr_rep3_R1_L002.fastq.gz
```

```
tp53_ko_rep1_R1_L001.fastq.gz
tp53_ko_rep1_R1_L002.fastq.gz
```

```
tp53_ko_rep2_R1_L001.fastq.gz
tp53_ko_rep2_R1_L002.fastq.gz
```

```
tp53_ko_rep3_R1_L001.fastq.gz
tp53_ko_rep3_R1_L002.fastq.gz
```

If your samples have been split across 3 lanes you are going to have 18 FASTQ files.

```
cntr_rep1_R1_L001.fastq.gz -
cntr_rep1_R1_L002.fastq.gz | cntr_rep1
cntr_rep1_R1_L003.fastq.gz -
```

```
cntr_rep2_R1_L001.fastq.gz -
cntr_rep2_R1_L002.fastq.gz | cntr_rep2
cntr_rep2_R1_L003.fastq.gz -
```

```
cntr_rep3_R1_L001.fastq.gz -
cntr_rep3_R1_L002.fastq.gz | cntr_rep2
cntr_rep2_R1_L003.fastq.gz -
```

```
tp53_ko_rep1_R1_L001.fastq.gz -
tp53_ko_rep1_R1_L002.fastq.gz | tp53_ko_rep1
tp53_ko_rep1_R1_L003.fastq.gz -
```

```
tp53_ko_rep2_R1_L001.fastq.gz -
```

```
tp53_ko_rep2_R1_L002.fastq.gz | tp53_ko_rep2
tp53_ko_rep2_R1_L003.fastq.gz -
```

```
tp53_ko_rep3_R1_L001.fastq.gz -
tp53_ko_rep3_R1_L002.fastq.gz | tp53_ko_rep3
tp53_ko_rep3_R1_L003.fastq.gz -
```

Note that when samples are split across lanes they become technical replicates and they are typically merged into single file. Also note that if you instead had paired end data you will have to further multiple by 2 to get total number of FASTQ file.

```
cntr_rep1_R1_L001.fastq.gz -
cntr_rep1_R2_L001.fastq.gz |
|
cntr_rep1_R1_L002.fastq.gz | cntr_rep1
cntr_rep1_R2_L002.fastq.gz |
|
cntr_rep1_R1_L003.fastq.gz |
cntr_rep1_R2_L003.fastq.gz -
```

```
cntr_rep2_R1_L001.fastq.gz -
cntr_rep2_R2_L001.fastq.gz |
|
cntr_rep2_R1_L002.fastq.gz | cntr_rep2
cntr_rep2_R2_L002.fastq.gz |
|
cntr_rep2_R1_L003.fastq.gz |
cntr_rep2_R2_L003.fastq.gz -
```

```
cntr_rep3_R1_L001.fastq.gz -
cntr_rep3_R2_L001.fastq.gz |
|
cntr_rep3_R1_L002.fastq.gz | cntr_rep2
cntr_rep3_R2_L002.fastq.gz |
|
cntr_rep2_R1_L003.fastq.gz |
cntr_rep2_R2_L003.fastq.gz -
```

```
tp53_ko_rep1_R1_L001.fastq.gz -
tp53_ko_rep1_R2_L001.fastq.gz |
|
tp53_ko_rep1_R1_L002.fastq.gz | tp53_ko_rep1
tp53_ko_rep1_R2_L002.fastq.gz |
|
tp53_ko_rep1_R1_L003.fastq.gz |
tp53_ko_rep1_R2_L003.fastq.gz -
```

```
tp53_ko_rep2_R1_L001.fastq.gz -
tp53_ko_rep2_R2_L001.fastq.gz |
|
tp53_ko_rep2_R1_L002.fastq.gz | tp53_ko_rep2
tp53_ko_rep2_R2_L002.fastq.gz |
|
tp53_ko_rep2_R1_L003.fastq.gz |
tp53_ko_rep2_R2_L003.fastq.gz -
```

```
tp53_ko_rep3_R1_L001.fastq.gz -
tp53_ko_rep3_R2_L001.fastq.gz |
|
```

```
tp53_ko_rep3_R1_L002.fastq.gz | tp53_ko_rep3
tp53_ko_rep3_R2_L002.fastq.gz |
|
tp53_ko_rep3_R1_L003.fastq.gz |
tp53_ko_rep3_R2_L003.fastq.gz -
```

5.3 Unix tooling

There are many different biology / bioinformatics related file formats. You will naturally come across different files formats. It doesn't matter so much what they are called, i.e what they file extension is, rather what they hold.

Essentially all bioinformatics files are either plain text (a.k.a ASCII) or binary. Plain text files are “easy” to interface with. Unix/Linux command line provide many great tools that can work with plain text files. To list a few popular and powerful command line tools:

- **grep** search for a pattern
- **sed** search and replace
- **awk** search and replace and more
- **cat** look inside
- **cut** what the knife would do to bread

It is out of the scope of this course to go through any of the tools or command line in general, but if you are going to do any kind of bioinformatics you'll eventually going to come across these tools.

The other file type is binary, this will requires designated tool for interfacing with the file for example SAM/BAM files have **samtools** that aid interaction with them.

- which coordinate on that chromosome read map to
- some quality metric on how well it mapped, all bases mapped or just some?

This is a bare minimum information you want your aligner to output, but there is actually a lot more information aligners typically output including was the read a multi-mapper, the orientation of read in terms of strand and position first or second, this is in the case of paired-end data. Some aligners like STAR¹ also include information about splice junctions. Also with sample and UMI barcoding there are moves to ask aligners to record that information as well.

Essentially all aligners these days will produce a sequence alignment map or sam file for short. Sam file is once again is a simple plain text file with rows and columns tab separated, with a header section that is unique to a sam file. It is best practice to keep sam files zipped, but instead of conventional `gzip` utility use special for this purpose `samtools` utility and the convention is to call compressed files BAM for binary sam. An example below illustrates relatively extreme case, never the less show a massive advantage in disk space in using bam files instead of sam.

```
-rw-r--r-- 1 kirill kirill 41G Aug 30 14:01 cntr_rep1_sorted_mdups.bam
-rw-rw-r-- 1 kirill kirill 399G Aug 31 11:06 cntr_rep1_sorted_mdups.sam
```

We always say that you should never store sam files on disk. Always convert them to bam files. Bam files are an industry standard, a lot of bioinformatics analysis, rotates around bam files. There even some thoughts to keep raw data in bam files instead of FASTQs.

One important thing you need to know about aligners is that broadly speaking there are two types, splice aware aligners and splice unaware aligners. This is different from gapped alignment. Essentially any good aligners will do gapped alignment to account for insertions and deletions of few bases. However in the case of RNAseq some gaps can actually be quite large and can be few kilobases, those are the splice junctions. Below is a list of a few aligners that are commonly used in our group, but of course there are many more aligners out there.

Don't use these ones for RNAseq analysis, these ones are for DNA mapping, including data from ChipSeq and ATACseq.

Common splice unaware aligners:

- `bowtie2`
- `bwa`

These aligners commonly used for RNAseq analysis. You can use for ChipSeq and ATACseq as well, but best not to, unless you know what you are doing. This is mainly because these aligners will model gaps differently, since large gaps (splice junctions) are expected you might get wrong scoring of insertion and deletions in your Chip and ATAC seqs.

Common splice aware aligners:

- `tophat2`
- `hisat2`
- `STAR`

Let's talk a little more about bam files, before moving on to read counting as a step two in typical RNAseq analysis.

6.2.1 Bam files

Bam files can be sorted in a couple of ways:

- by coordinate (4th column): arrange reads from first base till last base on each chromosome. Remember that chromosomes themselves may not be sorted by name or length
- by name (1st column): arrange reads alphabetically on each chromosome. Remember each read in a FASTQ file has a unique name (yes all 20 million or more reads).

The reason for name sorting is to place R1 and R2 reads next to each other (one after another), this mainly applicable for paired-end sequencing. I can't see any effect sorting by name single-end data, this

¹<https://github.com/alexdobin/STAR>

will make it as random as not sorting at all. Sorting by coordinates is a dominant way of sorting bam files, big advantage of sorting by coordinates is ability to for quick lookup of any region with in the bam. You will need index bam files.

Remember that you should end up with one bam file per sample and typically each bam will be also sorted. You don't need both copies of sorted and unsorted bam files. Both contain the same information. Sorted bam files will result in slightly reduced bam file sizes, this is because sorted bam files will give slightly better compression.

```
cntr_rep1_R1.fastq.gz  cntr_rep1.bam  cntr_rep1_sorted.bam
cntr_rep2_R1.fastq.gz  cntr_rep2.bam  cntr_rep2_sorted.bam
cntr_rep3_R1.fastq.gz  cntr_rep3.bam  cntr_rep3_sorted.bam
```

```
tp53_ko_rep1_R1.fastq.gz  tp53_ko_rep1.bam  tp53_ko_rep1_sorted.bam
tp53_ko_rep2_R1.fastq.gz  tp53_ko_rep2.bam  tp53_ko_rep2_sorted.bam
tp53_ko_rep3_R1.fastq.gz  tp53_ko_rep3.bam  tp53_ko_rep3_sorted.bam
```

6.3 Read counting

6.3.1 Introduction

Before leaping into what read counting is let's remind ourselves quickly what an RNAseq data looks like if we visualise it. There are number of different applications that you can use to visualise, but the one we most commonly use is Integrative Genomics Viewer (IGV)²

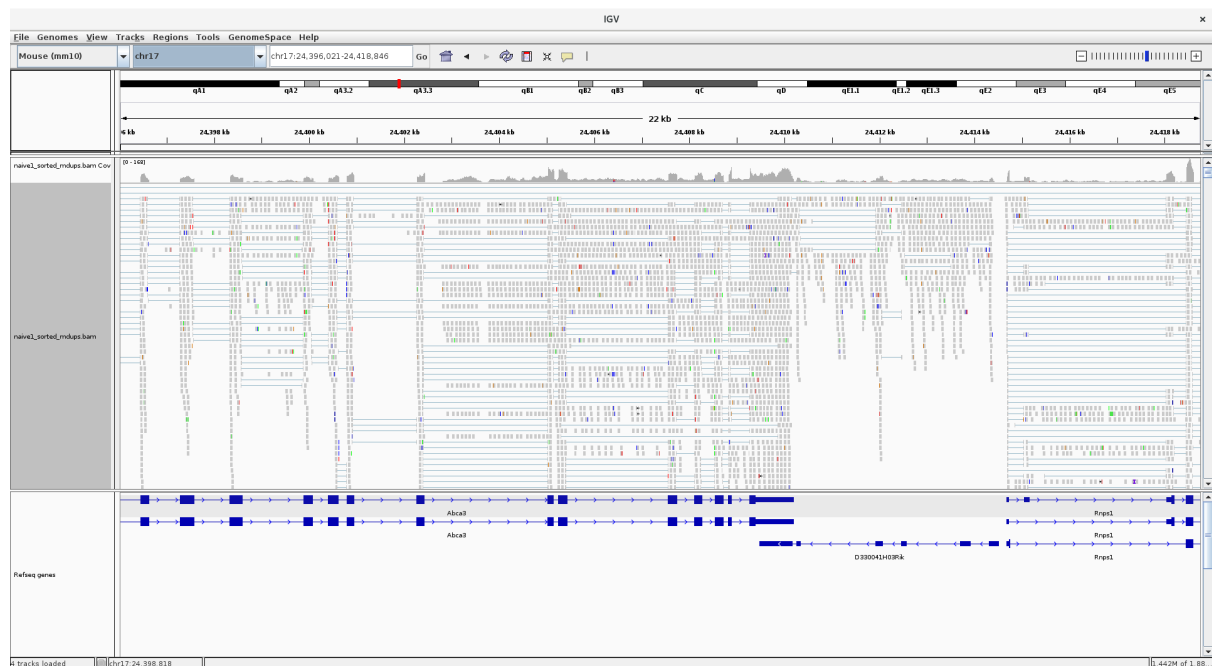
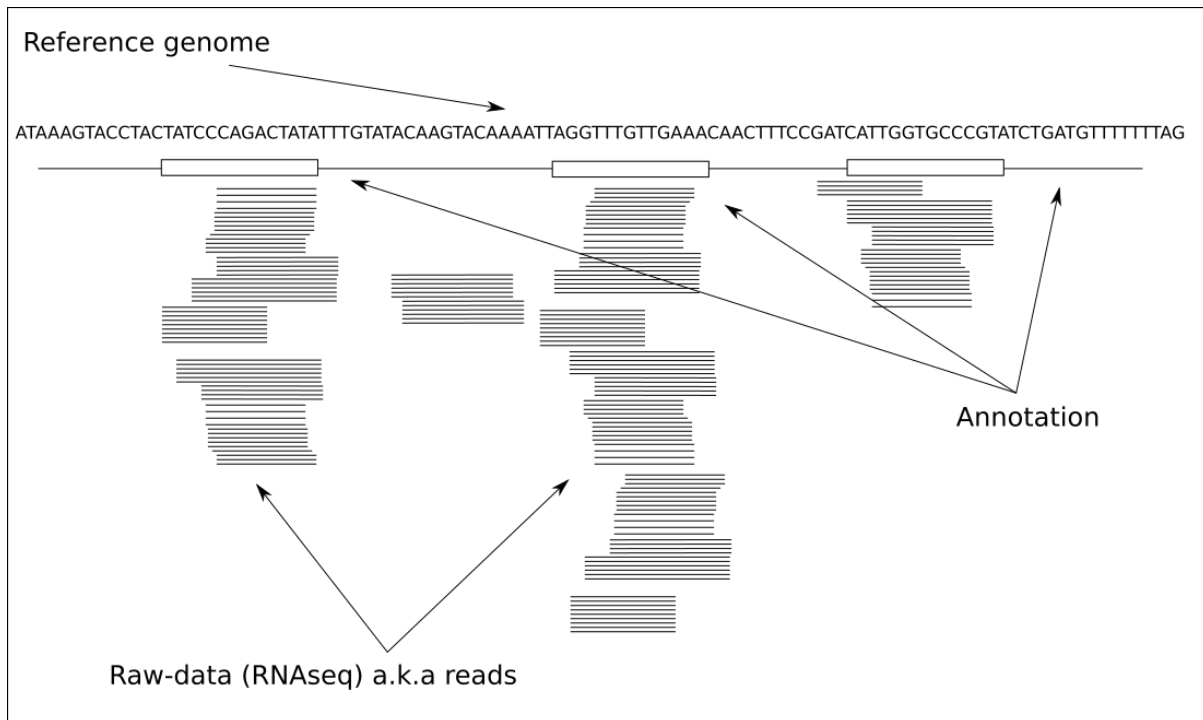


Figure 6.2: IGV browser view of one sample

Typically however there are multiple samples you'd be looking, which I left out from the image for clarity of view.

Here is schematic view of what you typical see in a genome viewer

²<http://software.broadinstitute.org/software/igv/>



The three things that we always need for visualisation of that kind are:

- reference genome file (FASTA)
- annotation file (GTF/GFF)
- data:
 - bam: these are information rich, but slow to load and navigate
 - bigWig: one example of coverage plot file, just show the coverage, small and quick
 - vcf: variant information, not part of this course

Note that something like IGV already comes with in built reference and annotation files, but just remember that it's likely comes with different version of the genome. Often time IGV reference files will be good enough for first look and if you start suspecting something you should load your specific reference.

Go through igv screen shoots and explain few different concepts, like overlapping features

Now that we've recapped what RNAseq data typically looks like lets talk how we approach read counting bioinformatically

6.3.2 Table of counts

Given an alignment map file a.k.a bam file and an annotation file e.g GTF we need to systematically go through every read in the bma file and assign it to a feature in the annotation file. Ultimately we would like to get contiguous table or simply, table of counts, where each row is a gene and columns are our samples. Here is a top few lines from typical table of counts.

Gene.ID	Chrom	Gene.Name	Biotype	K03_S3	K04_S4	WT1_S1	WT2_S2
ENSMUSG00000051951	1	Xkr4	protein_coding	0	0	0	0
ENSMUSG00000025900	1	Rp1	protein_coding	0	0	0	0
ENSMUSG000000109048	1	Rp1	protein_coding	0	0	0	0
ENSMUSG00000025902	1	Sox17	protein_coding	0	0	0	0
ENSMUSG00000033845	1	Mrpl15	protein_coding	934	1317	888	939
ENSMUSG00000025903	1	Lypla1	protein_coding	705	848	647	747
ENSMUSG000000104217	1	Gm37988	protein_coding	0	0	0	0
ENSMUSG00000033813	1	Tcea1	protein_coding	354	436	368	442
ENSMUSG00000002459	1	Rgs20	protein_coding	0	0	0	0

Note that this table has a few extra columns with some extra information, those aren't strictly necessary, but do allow to put things into biological picture.

This table would be minimal information you need for differential analysis. Note that `Gene.ID` column has to be unique values.

Gene.ID	K03_S3	K04_S4	WT1_S1	WT2_S2
ENSMUSG00000051951	0	0	0	0
ENSMUSG00000025900	0	0	0	0
ENSMUSG00000109048	0	0	0	0
ENSMUSG00000025902	0	0	0	0
ENSMUSG00000033845	934	1317	888	939
ENSMUSG00000025903	705	848	647	747
ENSMUSG00000104217	0	0	0	0
ENSMUSG00000033813	354	436	368	442
ENSMUSG00000002459	0	0	0	0

In RNAseq we typically should see reads only around feature, all reads should remain within the feature coordinates boundary. Reads should pile up on exons with clear boundaries between adjacent exons. Reads can and do map anywhere though and if reads fall outside of “normal” regions it could be an interesting case to follow up. As illustrated before visualisation is very important part of working with RNAseq data.

A few things to know when visualising RNAseq data:

- use the same reference sequence that you mapped against, otherwise you might see more variance between reference genome and the data
- be sure to visualise against the same annotation version and vendor type. This will ensure that your features (boxes) align best with your data. This is more important than picking the same reference files as features do change coordinates from version to version. Therefore if you are using different version of annotation to which you used for read assignment for visualisation you might see too many reads outside your feature boundaries, leading to false sense that there is some biological effect going on.
- typically RNAseq viewers sub-sample your data for visualisation purposes. This is mainly for speed and performance of the viewer. This is especially for highly expressed genes with deep read coverage. Be aware of that when doing your exploratory analysis. Often times you can turn sub-sampling off. Once again not seeing all of the data may result in misleading observations.

6.3.3 Details

We normally use words like “counting features” or “read assignment to features”, all means the same obtain table of counts. It is important to appreciate that “feature” could mean several things, it is some region of DNA that has some value/interest, a protein coding gene is a DNA region that produces a protein. For the purpose of RNAseq experiment we are most likely interested in protein coding features and therefore need to count those.

Similar to aligners there are number of tools out there that do read counting, but the two most popular are:

- `htseq-count`³
- `featureCounts`⁴

One thing to mention is that STAR aligner, mentioned previously, can actually output count files as well. In essence one can do most of RNAseq processing with STAR aligner.

I'm most familiar with how `featureCounts` works and it is a versatile tool, by default it counts reads that mapped to exons, which are referred as features, and summarises at meta-feature level which by default are gene. The number that ends up in the table of counts is a sum of all of reads from exons in a give gene. Remember that if a read overlaps more then one feature, by default, it is marked as an

³https://htseq.readthedocs.io/en/release_0.10.0/index.html

⁴<http://subread.sourceforge.net/>

ambiguous read and isn't counted toward a feature. It is possible to summarise reads at exon level which opens a possibility for differential exon usage analysis from RNAseq, something that isn't covered in this course.

We are going to through more QC's in the next section but here is nice summary image from `htseq-count` docs, to cement all possible read assignment problems.

	union	intersection_strict	intersection_nonempty
	gene_A	gene_A	gene_A
	gene_A	no_feature	gene_A
	gene_A	no_feature	gene_A
	gene_A	gene_A	gene_A
	gene_A	gene_A	gene_A
	ambiguous (both genes with --nonunique all)	gene_A	gene_A
	ambiguous (both genes with --nonunique all)		
	alignment_not_unique (both genes with --nonunique all)		

Figure 6.3: https://htseq.readthedocs.io/en/release_0.10.0/count.html

The final step in RNAseq analysis is to take table of counts and apply statistical modeling to study differential expressions of genes. This will be covered in later sections.

6.4 RNAseq checklist

Having introduced you to and RNAseq workflow below is a checklist for you to take home. It is broken into three sections, must, should and up to you, reflecting the type of information needed for the analysis. However remember that computer will always give you an answer based on the information you are going

to feed into it. It is in your best interests to provide as much information about your experiment as you can to the bioinformatician for most accurate analysis.

The short answer to a typical question by researcher to a bioinformatician; “What do you need to know about the experiment?” Everything ! but here is a break down of everything:

6.4.1 Must

If the below information isn't give I can't really do the analysis

- Have you got raw (FASTQ) files?
- What is your model organism?
- Which reference database to use? (e.g. ucsc, ensembl, refseq)
- Additional reference files e.g transgenes? will need relevant files (FASTA , GTF/GFF)
- Do you have a samples sheet for me?

6.4.2 Should

I'm saying should, but really the more information you provide the more accurate - more reflective analysis is going to be of your biological experiment

- Have I got enough explanation about experimental design including:
 - Any batch effects. Paired-data (eg. individual before/after treatment)
 - Comparisons of interest. Pairwise? Interaction?
 - Additional factors for contrast matrix:
 - * sample pairing
 - * sex
 - * time points
 - * phenotype
 - * stimulus
 - * other..

6.4.3 Up to you

This is your experiment and your money and time spend on it. I'm just saying that there are instrument to instrument variations and biases. As well as certain artifacts due to particular library preparation method. On top of that you'll need this information if you are going to publish your results in a paper anyway.

- Which sequencing facility was doing the sequencing?
- Library preparation info:
 - library type: single or paired end
 - preparation method ribo-depletion or poly(A), other
 - was library stranded
 - name of preparation kit was used
- Sequencer used (e.g HiSeq1500, NovaSeq, NextSeq etc)
- Are you keeping a copy of your raw data somewhere safe

Remember that the best practice is to gather all of the information in the checklist before the analysis

6.5 Summary

- RNAseq analysis is a relative measure of reads between samples.
- step one in RNAseq analysis is to get bam files
- step two in RNAseq analysis is to get table of counts.

- step three in RNAseq analysis do differential expression

Question: What is the correct order for an alignment based RNAseq workflow?

- differential expression, mapping, read counting
- read counting, mapping, differential expression
- mapping, read counting, differential expression
- differential expression, read counting, mapping

Chapter 7

QC

I hope that everyone appreciates important of checking things for quality. Any kind of work needs to be checked for some kind of integrity, being it agarose gel run, testing for whether the bands - DNA/RNA is present, or some spectrophotometer technique testing absorbs of light at 260nm. Similarly in bioinformatics world we also quality checks or QCs for short. I've broken this chapters into different QC's we typically do, dedicating a section for each type of QC.

7.1 Number of mapped reads

This would be one of the simplest checks that one should do or ask for, how many of my sequenced reads mapped to the reference genome? In previous chapter it was mentioned that the reference genome of any species is work in progress and that it'll will change with time and between different vendors. Although this is a become less true for mouse (*Mus_musculus*) genome as it has been extensively studied, but never the less one shouldn't assume the genome is static and is the "golden reference" so to speak.

Typically for mouse RNAseq data we tend to see anywhere between 80 and 99 percent. It is rare to see 100 % mapping to the reference in mouse at least and if you did I'd be a little worried. You should however not think of this or any other metrics as plain cutoff at a set value, rather think about what that means. Having seen a few mouse RNAseqs I'm not concern if mapping rate is above 80 %. If I'm seeing mapping rate lower than that. I have to start asking the question; Where do those 20-30% of reads come from?

- Is my reference not of high quality?
- Have the samples been contaminated?
- Why those reads don't map?
 - How long are the reads? (shorter reads harder to map)
 - skewed library?
 - bad quality reads (issue at the sequencer machine?)

As you can see there are number of possible things that could happened. One should brain storm all possible ways thing could have gone wrong and then try to rule out and thereby narrow down on possible problems. In general different QC metrics augment each other. This is why a tools like MultiQC¹ is a great way to aggregate result in a single report.

In the case of mouse genome, it is unlikely to be the problem with the genome quality, unless you've used some alternative version, so we can rule that out. As somebody who potentially extracted RNA you should be able to go back and check the lab book notes for wet-lab QC's and rule contamination out on your end. Perhaps sequencing facility didn't do a good job on they end.

¹<http://multiqc.info/>

The good place to start troubleshooting low mapping is to simply grab unmapped reads and BLAST (basic local alignment search tool)² against all possible species databases (NCBI). Provided the read(s) had a hit, this should give you pretty good understanding of the potential contamination source. The reads don't have to have a hit, and if they don't once again you need to think critically what that means. Is there something in your biology that could lead to unknown reads or chimeric reads or once again an artifact at the sequencing facility end where the machine just produces "garbage" reads.

FastQC³ is another excellent places to start checking your reads quality and overrepresented sequences, which once again should give a hit of why you reads don't map as well. We will talk about FastQC report later.

Having established what mapping rate is lets dig a little deeper into different types of "mapped reads", after all not all reads that mapped can be used for RNAseq analysis later on.

7.2 STAR aligner

There is really three types of reads as far as aligner concerned:

- mapped uniquely i.e a single location in the genome
- mapped not uniquely i.e mapped to multiple places in the genome
- not mapped at all

Multi-mapping reads usually split into a couple of classes. The best way to describe a multi-mapper is to give it some score - probability how likely that a read belongs to a particular location in the genome and bwa aligner does exactly that.

MAPQ: MAPping Quality. It equals $-10 \log_{10} \text{Pr}\{\text{mapping position is wrong}\}$, rounded to the nearest integer. A value 255 indicates that the mapping quality is not available.

I can't give you details as to why, but most other aligners use slightly different scheme, they basically assign an integer value instead:

- 50|255 = Uniquely mapping
- 3 = Maps to 2 locations in the target
- 2 = Maps to 3 locations in the target
- 1 = Maps to 4-9 locations in the target
- 0 = Maps to 10 or more locations in the target

Depending on the aligner, but there will be some threshold at which it will stop trying to map a multi-mapping read further. For STAR that threshold is set to 10 location, you can tweak that parameter to be what ever you like.

For STAR there are five categories of reads:

- uniquely mapped
- multi-mapped to less then 10 loci
- multi-mapped to more then 10 loci
- unmapped too short
- unmapped

Remember discussion about soft-clipping reads from previous section. If the read had been soft-clipped beyond the threshold then it'll be marked as too short and discarded. Going back to troubleshooting of unmapped reads if you are getting large proportion of reads that are too short, then despite what being said in the trimming section it is worth while trying to adapter and quality trim before mapping. Perhaps your library has high adapter content.

²<https://blast.ncbi.nlm.nih.gov/Blast.cgi>

³<https://www.bioinformatics.babraham.ac.uk/projects/fastqc/>

7.2.1 Details

There is a slight variation in how different tools count and interpret mapped reads. In terms of terminology there are three important concepts one need to understand:

- library size / total number of reads
- primary and secondary alignments
- supplementary alignment

Library size in this context means total number of reads in the FASTQ files, this is typically in order of millions of reads, 20 million reads per sample seems to be used as a rule of thumb for RNAseq analysis, but that number can be anywhere between 5 and 80 millions per sample. Remember from previous section that we are trying to map all reads from FASTQ to the reference. Therefore you expect to find all 20 million in your bam file. Bam file holds one alignment per line and remember that reads can have multiple alignments, there fore one read could occupy multiple lines. If the read maps to multiple location one of those location will be marked as primary (a.k.a primary alignment) all other will be marked as secondary. It is outside of the scope of this book how primary alignment gets picked, sometimes this is rather random process. Supplementary alignment on the other hand marking a read that has chimeric alignment such mapping between two chromosomes or an inversion of some kind. A uniquely mapped reads are the ones that actually going to make into differential analysis those are the read that have been marked as primary but also carry mapq of 255.

This command will give you the number of uniquely mapped reads that are good for differential expression

```
samtools view -c -F 256 -q 255 ctrl_rep1_sorted_mdups.bam
```

To actually get number of reads that mapped to the reference one need to keep track of number of unmapped reads as well and output them to bam, this is usually default behaviour. If you see 100% mapping rate I would suspect first thing that you didn't output unmapped reads into bam file

Also note that for paired end library total number of reads is double to what it is in the fastq file, because there is two fastq files R1 and R2. While STARs interpretation of reads isn't strictly correct it is actually the correct number of fragments that you need to consider for differential analysis.

7.3 Number of assigned reads

The second most important metric that you should look at is number of reads assigned to a feature. `featureCounts` provides nice summary metrics about read assignment. Remember just because your reads have mapped to the reference genome it doesn't mean your reads came from protein coding reads. We would like to see as many reads assigned as possible and once again there isn't hard threshold that one can apply to see if read assignment was "good" or "bad". Think about biology and what it means in the context of your experimental design.

Below are some of typical metrics for `featureCounts` that one should look at. There isn't enough time in this book to cover every metrics output by `featureCounts`. These metrics will be covered in details:

- Assigned: reads mapped to a feature
- Unassigned Ambiguity: overlapping with two or more features
- Unassigned MultiMapping: reads that map to more than one locations in the genome
- Unassigned NoFeatures: not overlapping with any features included in the annotation.
- Unassigned Unmapped: reads are reported as unmapped in SAM/BAM input. Note that if the '-primary' option of `featureCounts` program is specified, the read marked as

These metrics are for future reference only and are not discussed any further in the book:

- primary alignment will be considered for assigning to features.
- Unassigned MappingQuality: mapping quality scores lower than the specified threshold.
- Unassigned Chimera: two reads from the same pair are mapped to different chromosomes or have incorrect orientation.
- Unassigned Duplicate: reads marked as duplicate in the FLAG field in SAM/BAM input.
- Unassigned Secondary: reads marked as second alignment in the FLAG field in SAM/BAM input.

- Unassigned Nonjunction: reads that do not span exons will not be assigned if the ‘`-countSplitAlignmentsOnly`’ option is specified.
- Unassigned Overlapping Length: no features/meta-features were found to have the minimum required overlap length.

The one reads that go into differential analysis are assigned reads. All other reads are simply discarded for the purpose of RNAseq analysis. No feature reads depending on the study would be one of the most interesting “can of worms”. Together with other metrics no feature can suggest DNA contamination at it’s simplest. More interestingly though if DNA contamination isn’t true it is an interesting follow up of what those reads do and mean in your experiment, once again a good place to start is to BLAST them⁴. There isn’t much you can do with ambiguous reads. One can try a different annotation, but really the only work around would be to get longer reads and hope that they will be less ambiguous. Similar to ambiguous reads long reads could help with multi-mapping reads. One other possible solution is to look where reads multi-map and perhaps if reads only multi-map to one other location and that location hasn’t been annotated with a feature, reassign those reads to a known feature. This however is very speculative topic, which we won’t go into in this book. For our purpose unfortunately we simply going to discard those reads.

Coming back to a threshold value and what to expect. Typically for mouse RNAseq I tend to see approximately around 60% of reads being assigned to features. That is 60 % out of your 80 % or so mapped reads, so about 50 % of your total library goes into RNAseq analysis. It is a lossy process.

~_(-)~/~

7.4 FastQC

Overrepresented sequences could mean a couple of different things

- your library is highly duplicated
- issues at the library preparation
- issues at the instrument

7.5 Intragenic vs Intergenic region

7.6 Reads mapping bias

If the RNA wasn’t carefully treated and became fragmented then you are more likely to see 3 prime bias for reads mapping

7.7 Insert size distribution

7.8 Where to look for more help

- Biostarts
- SeqAnswers
- googling in general
- at the end of the day best way to get information is look at the source code OR ask the author

⁴<https://blast.ncbi.nlm.nih.gov/Blast.cgi>

Chapter 8

All about RNAseq

8.1 General

- Functional annotation of the mouse (FANTOM)
- Encyclopedia of DNA elements (ENCODE)
- Single loci can lead to to six different mRNA isoforms, on average, instead of single transcript

8.2 Technology

What is it for?

- what is expressed (exon locatios /strand) ?
- how much is expressed (hits per transcripts) ?

old way to do DE - microarray

- do need as much starting material for RNAseq

Noise:

- biological
- poisson noise from samplinig (have level of uncertanty at low level, low counts)
- technical
- Effect size (do the genes go up by alot or not - high signal)
- amount of variation

get better depth using poly-A enriched, because rRNA depletion will leave lots of other RNA around therefore diluting mRNA by some percent

mean fragment size ~ 300 bp

measure of uncertanty of the base call

8.3 Differential expression (DE)

- library size is total number of mapped reads. Library size might vary for individual samples.
- Total number of reads per gene proportional to:
 - gene expression level (to get cpm (count per million) divide each read count by library size multiplied by million). This is enable you to compare different library sized samples
 - transcript length

- sequencing depth of the library # Ideas
- have a url with multiqc report ready for browsing
- explain reference file (FASTA) and annotation files (GTF)

Chapter 9

Lesson plan

Perhaps it'll be good idea to introduce several different plain text files

- `txt`
- `csv/tsv`
- `gff/gff3/gtf`
- `fastq/fq`
- `fasta/fa/faa/fna`

DO NOT unzip your fastq file and store them as plain text ! DO NOT mess with your raw fastq files!

Chapter 10

Meeting notes

This is overall them of the day

How to instruct them to be good RNAseq user?

have a url with multiqc page ready?

This is

explain reference file (FASTA) and annotation files (GTF)

This if you can set up ftp link so that users can pont and click through directory structure to interact with the sikRun output.

Mentione salmon

Good to mention in parsing - salmon if you are using salmon you not going to see DNA contamination. so something to be aware.

Also mentione that this is for model organism that are well annotation/assembled.

Use Chiara's data to study chromosome 21. From multiqc report chromosome 21 has a lot more reads. It looks like ribosomal contamination even though library was polyA pooled.

Take all reads from chr21 and do feature counts and look where reads go. Suspecting that reads go to ribosomal genes. If so this would be great case study for the workshop.

<https://www.ncbi.nlm.nih.gov/pubmed/29788454>

genotify

Reproducible research using the same reference files

TCGA for example is using a single reference for the whole cohort

If you gonna do second study e.g chipSeq you gotta have it on the same reference as your RNAseq there is no such thing as "the reference genome" there are different builds, it is work in progress