

So you want to do a: Bulk RNA-Seq Differential Expression Analysis

Matthew L. Settles, PhD

Center for Advanced Technology

University of California, San Francisco

matt.settles@ucsf.edu; catbioinformatics@ucsf.edu

Disclaimer

- This talk/workshop is full of opinion, there are as many different way to perform analysis as there are Bioinformaticians.
- My opinion is based on over a decade of experience and spending a considerable amount of time to understand data, how its created and how it relates to the biological questions of interest.
- Each experiment is unique, this lecture is a starting place and should be adapted to the specific characteristics of your experiment.

What is Differential Expression

Differential expression analysis means taking *normalized* sequencing fragment count data and performing statistical analysis to discover *quantitative* changes in expression levels between experimental groups.

For example, we use statistical testing to decide whether, for a given gene, an observed difference in fragment counts between group A and group B is significant, that is, whether it is greater than what would be expected just due to natural random variation.

Goals

- Focus on **bulk** differential **gene** expression studies
- Discussion of considerations from sample collection through analysis and analysis output.

Treating Bioinformatics as a Data Science

Seven stages to data science

1. Define the question of interest
2. Get the data
3. Clean the data
4. Explore the data
5. Fit statistical models
6. Communicate the results
7. Make your analysis reproducible

**Data science done well looks easy
and that's a big problem for data
scientists**

**simplystatistics.org
March 3, 2015 by Jeff Leek**

Designing Experiments

Beginning with the question of interest (and work backwards)

- The final step of a DE analysis is the application of a statistical model to each gene in your dataset.

Traditional statistical considerations and basic principals of statistical design of experiments apply.

- **Control** for effects of outside variables, avoid/consider possible biases, avoid confounding variables in sample preparation.
 - **Randomization** of samples, plots, etc.
 - **Replication** is essential (triplicates are THE minimum)
-
- You should know your final (DE) model and comparison contrasts before beginning your experiment.

Power Analysis

- A systematic search resulted in six open-source tools for sample size calculation for RNA-seq differential expression analysis (**As of Jan 2017**)
- Exemplary sample size estimation performed by the six tools using real mouse and human data as input files led to widely different results.
- Tool evaluation using simulations showed that most tools estimate the sample sizes incorrect for low true effects and large dispersion.
- **Instead**, the use of pilot data is recommended.

Feasibility of sample size calculation for RNA-seq studies

Alicia Poplawski ✉, Harald Binder

Briefings in Bioinformatics, bbw144, <https://doi.org/10.1093/bib/bbw144>

Published: 18 January 2017 **Article history** ▼

Wu H, Wang C, Wu Z. PROPER: comprehensive power evaluation for differential expression using RNA-seq. *Bioinformatics* 2015;**31**:233–241.

Three outcomes

Goldilocks and the three bears

- Technical and/or biological variation exceeds that of experimental variation, results in 0 differentially expressed genes
- Experiment induces a significant phenotype with cascading effects and/or little to no biological variation between replicates (ala cell lines), results in 1000s of DE genes. Some of which are directly due to experiment; however, most due to cascading effects.
- Technical artifacts are controlled. Biological variation is induced in the experiment, and cascading effects are controlled, or accounted for, results in 100s of DE genes directly applicable to the question of interest.

General rules for preparing and experiment/ samples

1. Prepare more samples than you are going to need, i.e., expect some will be of poor quality, or fail
2. Preparation stages should occur across all samples at the same time (or as close as possible) and by the same person
3. Spend time practicing a new technique to produce the highest quality product you can, reliably
4. Quality should be established using fragment analysis traces (pseudo-gel images)
5. DNA/RNA should not be degraded
 1. RNA RIN > 7.0 is best
 2. 260/280 ratios for RNA should be approximately 2.0 and 260/230 should be between 2.0 and 2.2. Values over 1.8 are acceptable
6. Quantity should be determined with a fluorometer, such as a Qubit.

Sample preparation

In high throughput biological work (Microarrays, Sequencing, HT Genotyping, etc.), what may seem like small technical details introduced during sample extraction/preparation can lead to **large changes, or technical bias**, in the data.

Not to say this doesn't occur with smaller scale analysis such as Sanger sequencing or qRT-PCR, but they do become more apparent (seen on a global scale) and may cause significant issues during analysis.

BE CONSISTENT ACROSS ALL SAMPLES!!!

Generating RNA-Seq libraries

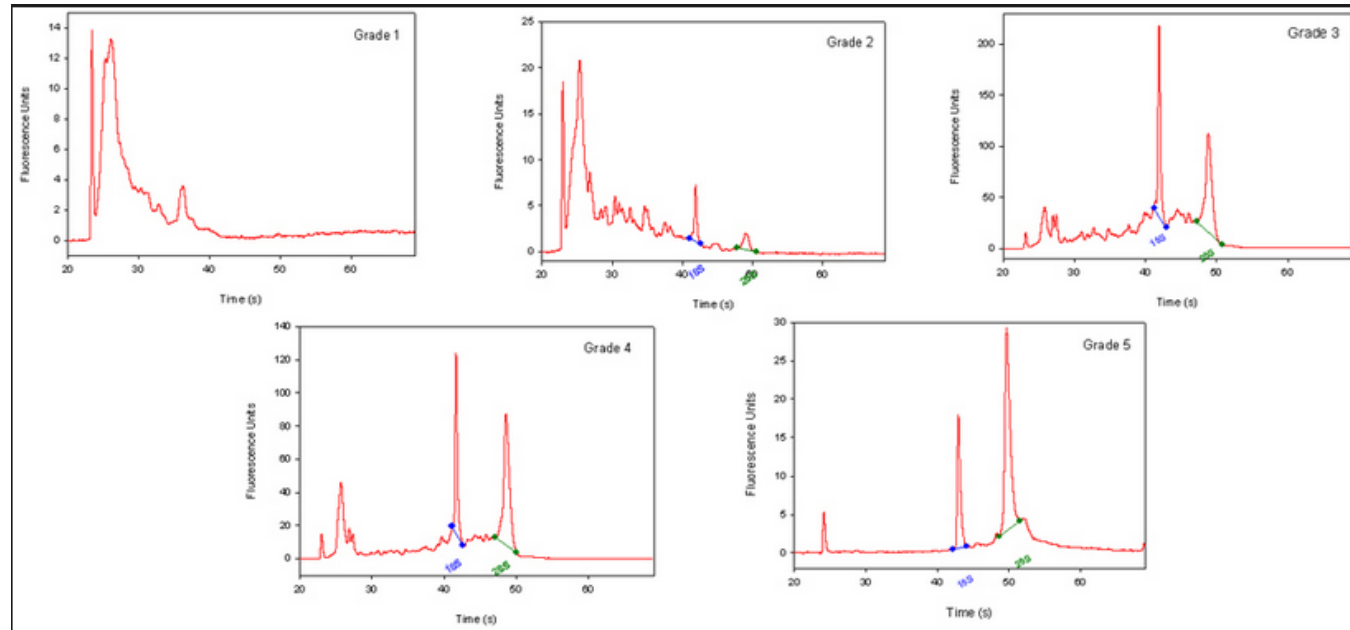
Considerations

- QA/QC of RNA samples
- What is the RNA of interest
- Library Preparation
 - Stranded Vs. Unstranded
 - Whole transcript Vs. 3-prime biased (tag-seq)
- Size Selection/Cleanup
 - Final QA

QA/QC of RNA samples

RNA Quality and RIN (RQN on AATI Fragment Analyzer)

- RNA sequencing begins with high-quality total RNA, only an Agilent Bioanalyzer (or equivalent) can adequately determine the quality of total RNA samples. RIN values between 7 and 10 are desirable.



BE
CONSISTANT!!!

RNA of interest

From “total RNA” we extract “RNA of interest”. Primary goal is to NOT sequence 90% (or more) ribosomal RNAs, which are the most abundant RNAs in the typical RNA sample. there are two main strategies for enriching your sample for “RNA of interest”.

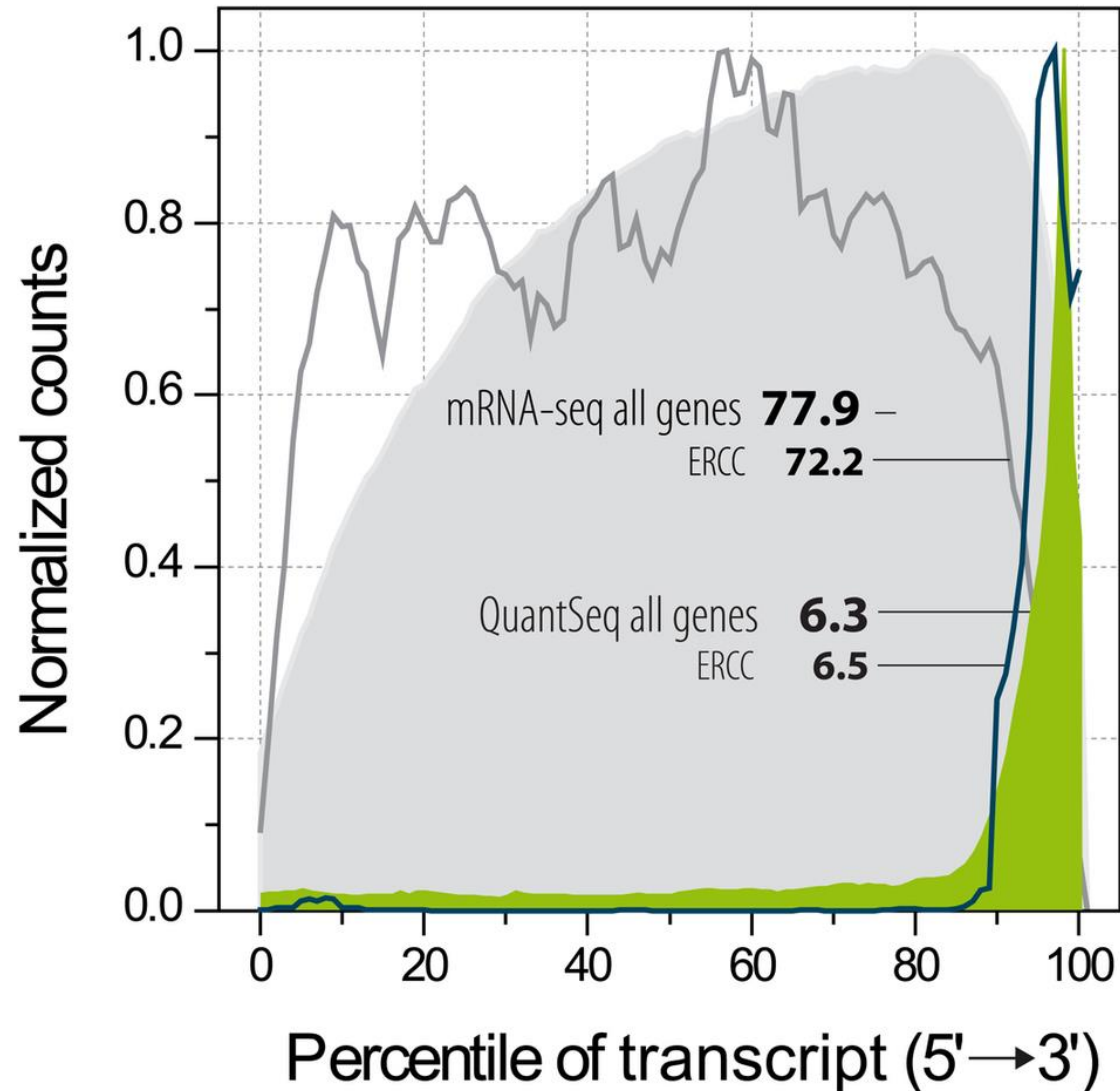
- polyA selection. Enrich mRNA (those with polyA tails) from the sample by oligo dT affinity.
- rRNA depletion. rRNA knockdown is mainly used when your experiment calls for sequencing non-polyA RNA transcripts and non-coding RNA (ncRNA) populations. This method is also usually more costly.

rRNA depletion will result in a much larger proportion of reads which align to intergenic and intronic regions of the genome.

Library Preparation

- Some library prep methods first require you to generate cDNA, to ligate on the Illumina barcodes and adapters.
 - cDNA generation using oligo dT (3' biased transcripts)
 - cDNA generation using random hexomers (less biased)
 - full-length cDNA generation will improve transcript and isoform analysis
- Also, can generate strand specific libraries, which means you only sequence the strand that was transcribed.
 - This is most performed using dUDP rather than dNTPs in cDNA generation and digesting the “RNA” strand.
 - Can also use an RNA ligase to attach adapters and then PCR the second strand and remainder of adapters.

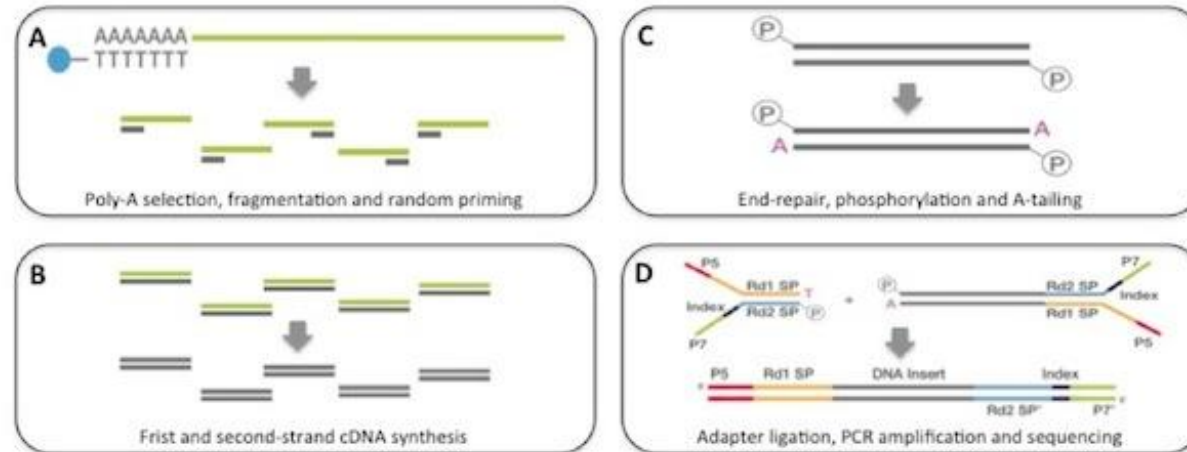
Library Preparation whole transcript vs 3' biased



Lexogen
[QuantSeq 3'](#)
[mRNA-Seq](#)

Protocol – whole transcript vs 3-prime biased

Illumina Tru-Seq RNA-seq protocol



Library prep begins from 100ng-1ug of Total RNA which is poly-A selected (A) with magnetic beads. Double-stranded cDNA (B) is phosphorylated and A-tailed (C) ready for adapter ligation. The library is PCR amplified (D) ready for clustering and sequencing.


Lexogen Quantseq protocol

Reverse transcription
(oligo-dT priming)

RNA removal

Second-strand
synthesis

Purification

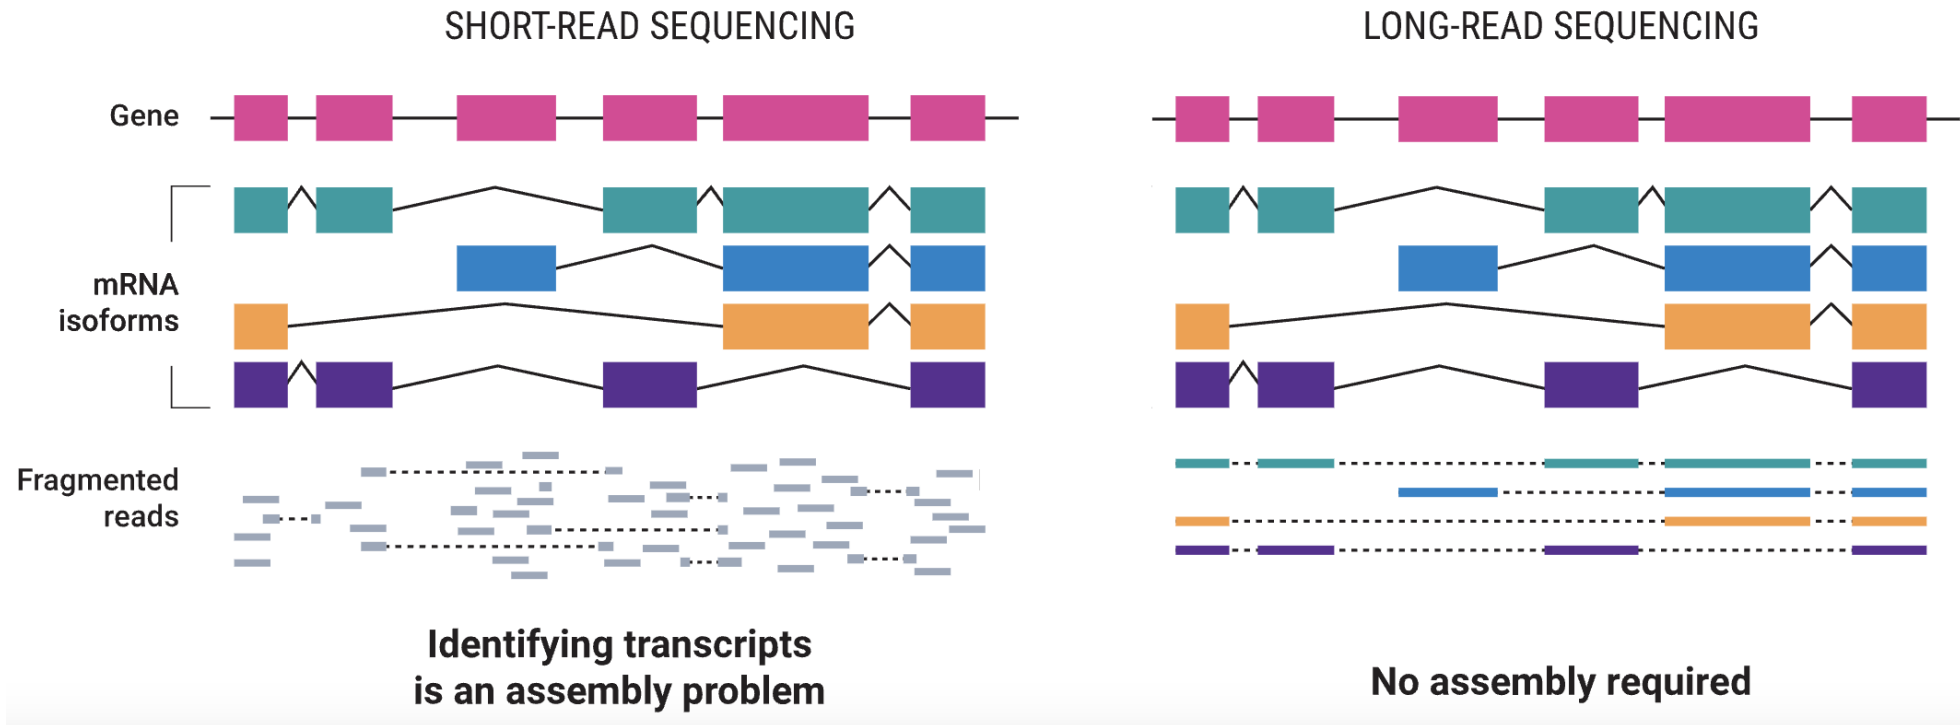


3-prime differences from whole transcript

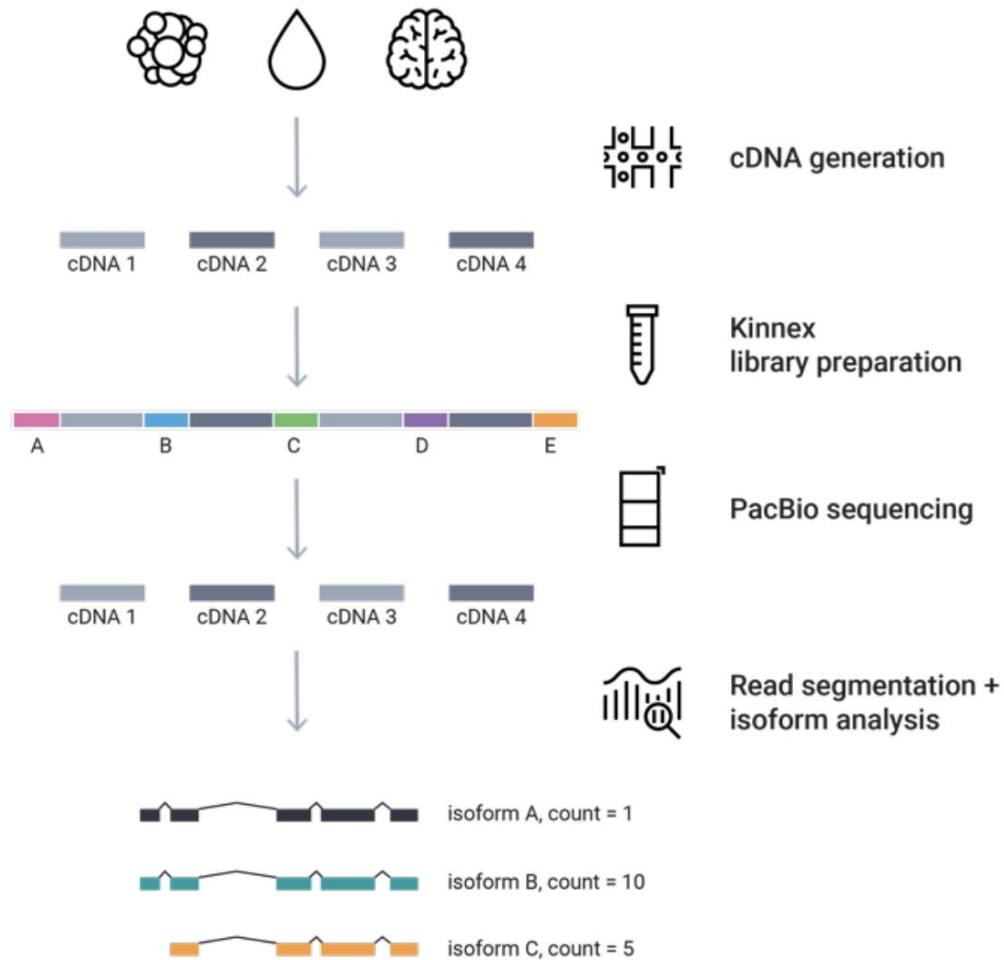
- Low input and low-quality samples (say as low as 100ng), **BUT protocol is a (a bit) more sensitive to chemical contaminants (spin column cleaned RNA samples are recommended)**
- Significant cost saving through faster library preparation protocol and multiplexing [need less sequencing per sample] comparable to microarrays.
- Best suited for gene counting as data **does not contain transcript splicing information on the whole transcript.**
- **Requires (sort of) a reference genome with good annotation (plus known UTRs)**
- **Only applicable to Eukaryotic samples (requires polyA)**
- Strand specific
- Single-end read sequencing is sufficient

PacBio - The future of RNA sequencing is with long reads!

ISOFORM DISCOVERY

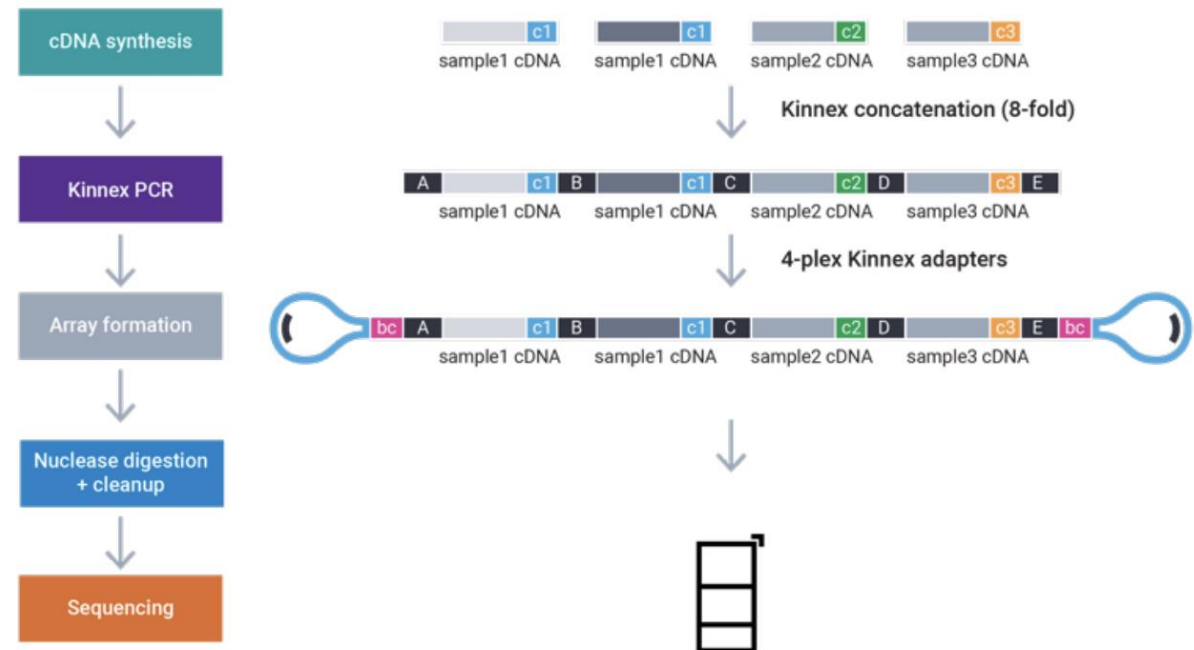


Full Length RNA + Kinnex + PB Revio Hifi



HiFi sequencing advantages for full-length RNA sequencing

- Sequence full-length isoforms from 5' to 3' ends
- Accurately characterize splice sites
- Detect novel genes and isoforms
- Obtain isoform read count information.
- With Kinnex increase throughput 12X and over 40M reads



Size Selection/Cleanup/QA

Final insert size optimal for DE are ~ 150bp (or kit suggestion)

- Very important to be consistent across all samples in an experiment on how you size select your final libraries. You can size select by:
 - Fragmenting your RNA, prior to cDNA generation.
 - Chemically heat w/magnesium
 - Mechanically (ex. ultra-sonicator)
- Cleanup/Size select after library generation using SPRI beads or gel cut
- QA the samples using a fragment analyzer (Bioanalyzer) and quantify with qPCR.

Most important thing is to be consistent!!!

Sequencing Depth

Coverage is determined differently for "Counting" based experiments (RNA-Seq, amplicons, etc.) where an expected number of reads per sample is typically more suitable.

The first and most basic question is how many reads per sample will I get
Factors to consider are (per lane):

1. Number of reads being sequenced
2. Number of samples being sequenced
3. Expected percentage of usable data
4. Number of lanes being sequenced

$$\frac{\text{reads}}{\text{sample}} = \frac{\text{reads.sequenced} * 0.8}{\text{samples.pooled}} \times \text{num.lanes}$$

Read length, or SE vs PE, does not factor into sequencing depth.

Sequencing

Characterization of transcripts, or differential gene expression

Factors to consider are:

- Read length needed depends on likelihood of mapping uniqueness, but generally longer is better and paired-end is better than single-end.
 - (2 x >100bp is best)
- Interest in measuring genes expressed at low levels (<< level, the >> the depth and necessary complexity of library)
- The fold change you want to be able to detect (< fold change more replicates, more depth)
- Detection of novel transcripts, or quantification of isoforms requires >> sequencing depth

The amount of sequencing needed for a given sample/experiment is determined by the goals of the experiment and the nature of the RNA sample.

Barcodes and Pooling samples for sequencing

- Best to have as many barcodes as there are samples
 - Can purchase barcodes from vendor, generate them yourself and purchase from IDTdna (example), or consult with the CAT core.
- Best to pool all samples into one large pool, then sequence multiple lanes
- IF you cannot generate enough barcodes, or pool into one large pool, RANDOMIZE samples into pools.
 - CAT core bioinformatics can produce a randomization scheme for you.
 - This must be considered/determined PRIOR to library preparation

[SUMMARY] Generating RNA-seq libraries

Considerations

- QA/QC of RNA samples [Consistency across samples is most important.]
- What is the RNA of interest [polyA extraction is recommended.]
- Library Preparation
 - Stranded Vs. Unstranded [Standard stranded library kits]
 - Whole transcript vs 3-prime biased [IF DE only, 3-prime biased]
- Size Selection/Cleanup [Target mean 150bp or kit recommendation]
 - Final QA [Consistency across samples is most important.]

Cost Estimation

- RNA extraction and QA/QC (Per sample)
- Enrichment of RNA of interest + library preparation (Per sample)
 - Library QA/QC (Bioanalyzer and Qubit)
 - Pooling [If you generate your own libraries]
- Sequencing (Number of lanes)
- Bioinformatics - General rule is to estimate the same amount as data generation, i.e. double your budget -

<https://cat.ucsf.edu/sequencing/sequencing-rates>

Illumina sequencing costs

NovaSeq X Series				
Flow cell type		1.5B	10B	25B
Meaningless	Single-end reads	1.6 billion	10 billion	26 billion
	Paired-end reads	3.2 billion	20 billion	52 billion

<https://www.illumina.com/systems/sequencing-platforms/novaseq-x-plus/specifications.html>

HOWEVER!!! The number of lanes per flow-cell are missing!

1.5B – 2 lanes

10B – 8 lanes

25B – 8 lanes

<https://cat.ucsf.edu/sequencing>

and

<https://cat.ucsf.edu/sequencing/sequencing-rates>

For additional help