



Summer Institute
in Statistical Genetics 2019

1. Experimental Design and Hypothesis Testing



ggibson.gt@gmail.com

<http://www.cig.gatech.edu>

Georgia
Tech



Course Outline

1. Experimental Design and Hypothesis Testing (GG)
2. Bulk Normalization (GG)
3. Single Cell RNAseq Basics (GG)
4. Single Cell Workflow (JP)
5. Single Cell Clustering (JP)
6. Single Cell Trajectory Analysis (JP)
7. Multiplexing; Tutorial / Discussion Session (JP)
8. Integrative Genomics (GG)
9. Heritability of Gene Expression (JP)
10. eQTL analysis (GG)

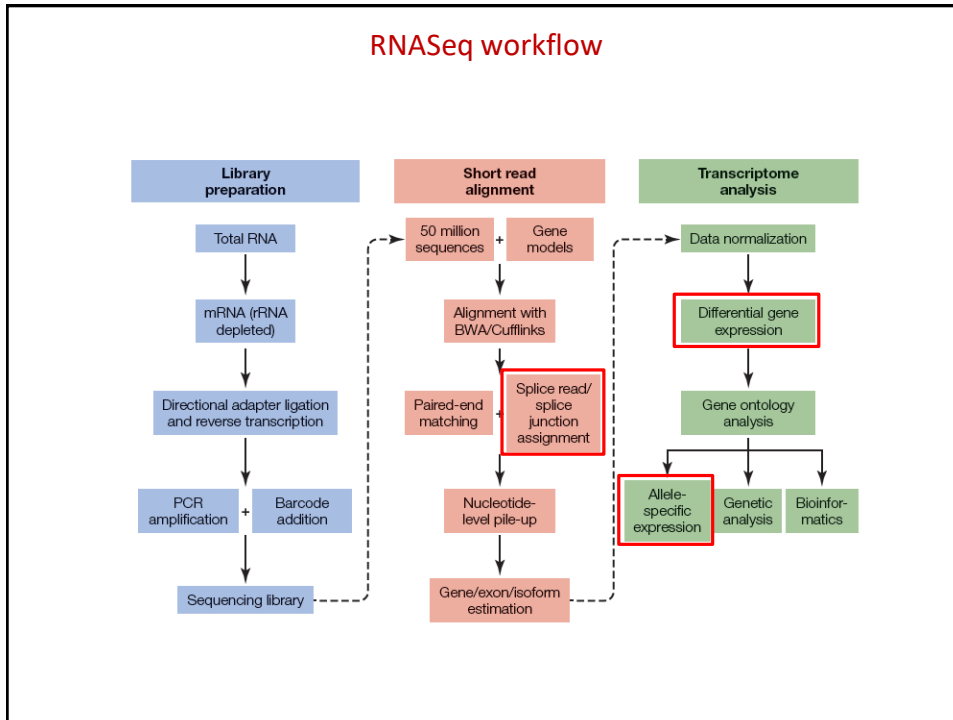
Steps in a Gene Expression Profiling Study

1. Experimental Design (this morning)
2. RNA Sequencing (next)
3. Short read alignment
4. Normalization (after the break)
5. Hypothesis testing (this morning)
6. Downstream analyses (module 10)
7. Genetic analysis (Friday afternoon)

RNA Sequencing

RNA is prepared, mRNA is captured on polyT beads, fragmented, and converted to cDNA using either a stranded or unstranded protocol, usually with 12-24X multiplexing

1. Single-end reads
 - Maximizes the total number of independent reads (50M optimal)
 - When RNA is degraded, eg FFPE specimens
2. Paired-end reads
 - Slightly more accurate alignment
 - But typically lower coverage (25M reads)
 - Better for estimation of alternate splicing and ASE
3. 3' targeted
 - Lexogen protocol is one fifth the cost (\$70 vs \$350 per sample)
 - Ideal for large sample studies when funds are a concern
 - Single Cell drop digital dd-scRNASeq is also 3' targeted



Core Open Source Software

While the Tuxedo protocol (Bowtie, Tophat, Cufflinks, Cuffdiff, CummeRbund) remains popular, we recommend the following open source alternatives.

1. Short Read Alignment

- STAR <https://github.com/alexdobin/STAR/releases>
- HISAT2 <https://ccb.jhu.edu/software/hisat2/index.shtml>

2. Read counting

- HTSeq <http://www-huber.embl.de/HTSeq/doc/overview.html>

3. Differential Expression

- DESeq <https://bioconductor.org/packages/release/bioc/html/DESeq2.html>
- DEXSeq <https://www.bioconductor.org/packages/release/bioc/html/DEXSeq.html>
- edgeR <https://bioconductor.org/packages/release/bioc/html/edgeR.html>
- Voom http://web.mit.edu/~r/current/arch/i386_linux26/lib/R/library/limma/html/voom.html

4. Data Normalization

- SVASEq <https://www.bioconductor.org/packages/release/bioc/html/sva.html>
- Combat <https://www.rdocumentation.org/packages/sva/versions/3.20.0/topics/ComBat>
- PEER <http://www.sanger.ac.uk/science/tools/peer>
- SNM <https://www.bioconductor.org/packages/release/bioc/html/snm.html>

Basic Experimental Design: Contrasts

At the same time, you need to be aware of the contrasts you wish to make since by tweaking the design you may gain a lot in terms of what you can infer.

Suppose you want to compare B cells and T cells from Healthy controls and Flu patients, and you have the funds to generate 24 RNASeq profiles



What is the best design?

- 6 controls and 6 patients, each donating both a B and a T cell sample
- 12 controls and 12 patients, each donating either a B or a T cell sample
- 3 controls and 3 patients, each donating a B and a T cell sample, processed twice
- 3 controls and 3 patients, each donating 2 B and 2 T cell samples, on separate days
- same as above, but only men or only women
- 12 controls and 12 patients, each donating either a B or a T cell sample, but pooling two visits

Main effects can only be contrasted if you have biological replicates:
reducing the number of individuals may allow you to address intra-individual variability

Interaction effects allow you to ask questions like whether B cells and T cells differ more between healthy volunteers or sick patients

Two hypothetical sets of results



Confounding Design Biases

At the design step, avoid confounding biological factors:

- don't contrast bloods from young males and old females
 - don't contrast hearts from normal mice and livers from obese ones
- as far as possible, balance all biological factors*

Be aware of the potential for technical confounding:

- date of RNA extraction or hybridization
- batch of samples (particularly for microarray studies)
- person who prepared the libraries
- SE or PE, read length and quality of reads
- quality of RNA (RIN = Bioanalyzer RNA Integrity Number)

Linear Modeling

1. Normalize the samples:

$$\log(\text{fluorescence}) = \mu + \text{Array} + \text{Residual}$$

OR variance transforms, OR supervised methods

2. For each gene, assess significance of treatment effects on the Residual (ie. Expression level) with a linear model:

$$\text{Residual} = \mu + \text{Sex} + \text{Geno} + \text{Treat} + \text{Interact} + \text{Error}$$

Hypothesis testing


1. Generally we are interested in asking whether there is a significant difference between two or more treatment group(s) on a gene-by-gene basis
2. For a simple contrast, we can use a t-test to test the hypothesis. Significance is always a function of:
 1. The difference between the two groups: [5,6,4] vs [7,5,6] has a diff of 1
 2. The variance within the groups: [2,5,8] vs [3,6,9] does as well, but is less obvious
 3. The sample size: [5,6,4,4,6,5] vs [7,5,6,5,6,7] is better
3. For contrasts involving multiple effects, we usually use **General Linear Models** in the ANOVA framework (analysis of variance: significance is assessed as the F ratio or between sample to residual sample variance.
4. **edgeR** uses **limma** to perform One-Way ANOVAs. This likelihood framework is very powerful, but constrains you to contrast treatments with a reference
5. Robust statistics (eg using **lme4**) also allow you to evaluate INTERACTION EFFECTS, namely not just whether two treatments are individual significant, but also whether one depends on the other
6. Given a list of p-values and DE estimates, we need to evaluate a significance threshold, which is usually done using False Discovery Rate (FDR) criteria, either B-H or a qvalue

T-tests and F-ratios


A **two-sample t-test** evaluates whether the two population means differ one another, given the pooled standard deviations of the sample and the number of observations:

$$t = \frac{\bar{X}_1 - \bar{X}_2}{s_p \sqrt{2/n}} \quad s_p = \sqrt{\frac{s_{X_1}^2 + s_{X_2}^2}{2}}$$

F-statistics generalize this approach by asking whether the deviations between samples are large relative to the deviations within samples. The larger the difference between the means, the larger the **F-ratio**, hence significance is evaluated by contrasting variances.



Generally we start by evaluating whether there is variation among all the groups, and then test specific post-hoc contrasts. With multifactorial designs you also should be aware of the difference between Type I and Type III sums of squares, namely univariate and conditional tests.



Hypothesis testing in edgeR

If you compare the output of edgeR to those of a t-test or F-test, particularly for small n, you will often get very different results. There are at least 4 reasons for this.

1. edgeR performs a **TMM normalization**. Since RNASeq generates counts, we adjust for library size by computing cpm (counts per million). If a few high abundance transcripts vary by several percent, they throw off all the other estimates. TMM fixes this.
2. edgeR shrinks the variance of low abundance transcripts by fitting the distribution to the “**negative binomial**” expectation. Basically it adds values to account for sampling error at the low end so that comparing 0, 1 and 2 is more like comparing 10, 11 and 12.
3. edgeR also employs a powerful **within-sample variance adjustment** in its GLM fitting, with the result that it puts much more weight on fold-change than standard F-tests.
4. For a one-way ANOVA the approaches are similar (though you need to be careful about whether you fit an intercept, which means you compare multiple samples to a reference rather than to one another). For more complicated analyses involving two or more factors, nesting, and random effects, the modeling frameworks are really quite different.

I prefer to run ProcGLM or ProcMIXED in SAS or JMP-Genomics, but there is no clear equivalent in R – lme4 is close, but needs looping.

GEO and ArrayExpress

The screenshot shows the NCBI GEO website. The main navigation bar includes 'GEO' and 'Gene Expression Omnibus'. Below the navigation, there are search filters for 'DataSets', 'Gene profiles', 'GEO accessions', and 'GEO BLAST'. A 'Submitter login' section is visible at the bottom left, with fields for 'User ID' and 'Password'. On the right side, there is a 'Public data' section with statistics: Platforms (7,411), Samples (479,499), and Series (17,145).

The screenshot shows the ArrayExpress website. The main heading is 'ARRAYEXPRESS'. Below it, there is a search bar and a 'Gene Expression Atlas' section. The 'Experiments Archive' section shows 11,755 experiments and 32,977 assays. There is a 'News' section with recent updates and a 'Links' section with various resources.