


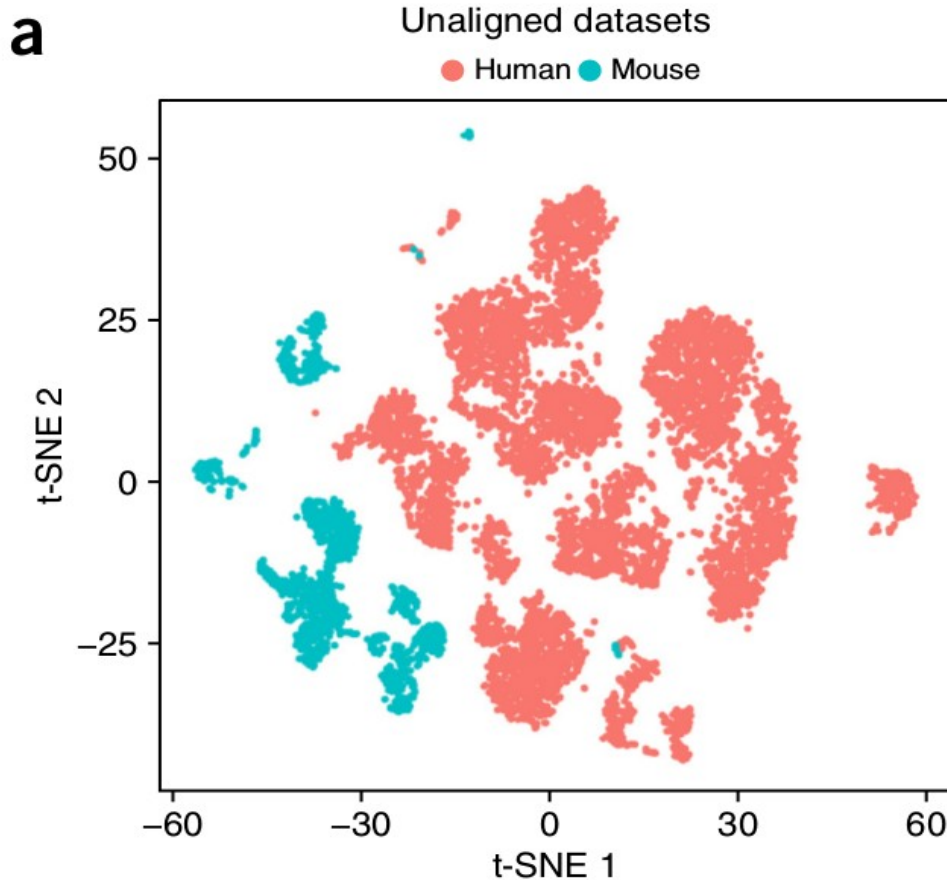
# Removing Technical Variation in scRNAseq Data

Nikolay Oskolkov  
NBIS Long-Term Support (WABI)



# Why to remove technical variation?

In order to facilitate discovering biological signal



## Batch-effects:

- 1) dates of sequencing
- 2) people done sequencing
- 3) flow-cells / plates
- 4) chemistry / protocol
- 5) lanes
- 6) read length
- 7) labs produced data
- 8) organisms
- 9) etc.

100% confounding: put cases  
and  
controls on different flow-cells


**Normalization:** correct for systematic variation in sequencing experiment

- 1) between samples (e.g. sequencing depth bias)
- 2) between features (e.g. gene length or GC content)

# How to detect technical variation?


## Difference in sequencing depth:

Sequencing Depth Across Cohorts



## Genome-Wide Batch-effects:


PCA PLOT, DONOR EFFECT



## Genome-Wide Batch-effects:


## Per-Gene Batch-effects:

Color Key




ILC scRNAseq

Adjusted R<sup>2</sup> of Association between PCs and Phenotypes



Observed vs. Resampled Variance Explained by Batch





# Combat Batch Effects Correction

Nikolay Oskolkov

May 7, 2018

## Generate Gene Expression with Batch

```
expr<-matrix(c(200,201,180,85,99,87,150,180,190,81,93,89),
             ncol=6,byrow=TRUE)
colnames(expr)<-paste0("S",seq(1:dim(expr)[2]))
rownames(expr)<-paste0("G",seq(1:dim(expr)[1]))
expr
```

```
##      S1  S2  S3 S4 S5 S6
## G1 200 201 180 85 99 87
## G2 150 180 190 81 93 89
```


Batch coding:

```
batch<-c(1,1,1,0,0,0)
batch
```

```
## [1] 1 1 1 0 0 0
```

## Visualize the Batch

```
plot(expr[1,]~c(1,1,1,0,0,0),col="blue",xlab="Batch",ylab="Gene Expr")  
abline(lm(expr[1,]~c(1,1,1,0,0,0)),col="blue")  
points(expr[2,]~c(1,1,1,0,0,0),col="red")  
abline(lm(expr[2,]~c(1,1,1,0,0,0)),col="red")
```



# Why not to Regress Out?

```
summary(lm(expr[1,]~batch))
```

```
##
## Call:
## lm(formula = expr[1, ] ~ batch)
##
## Residuals:
##      S1      S2      S3      S4      S5      S6
##  6.333  7.333 -13.667 -5.333  8.667 -3.333
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)    90.333     5.740   15.74 9.52e-05 ***
## batch          103.333     8.117   12.73 0.000219 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 9.941 on 4 degrees of freedom
## Multiple R-squared:  0.9759, Adjusted R-squared:  0.9699
## F-statistic: 162.1 on 1 and 4 DF, p-value: 0.0002194
```



## Why not to Regress Out?

We can just work with residuals:

```
residuals(lm(expr[1,] ~ batch))
```

```
##           S1           S2           S3           S4           S5           S6
##  6.333333  7.333333 -13.666667  -5.333333  8.666667  -3.333333
```

which is equivalent to mean-centering within each batch:

```
c(c(200,201,180)-mean(c(200,201,180)),c(85,99,87)-mean(c(85,99,87)))
```

```
## [1]  6.333333  7.333333 -13.666667  -5.333333  8.666667  -3.333333
```

## Correct with Limma

```
library("limma")  
removeBatchEffect(expr, batch)
```

```
##           S1           S2           S3           S4           S5           S6  
## G1 148.3333 149.3333 128.3333 136.6667 150.6667 138.6667  
## G2 107.1667 137.1667 147.1667 123.8333 135.8333 131.8333
```

Why is the result of "removeBatchEffect" so different from the residuals?

```
residuals(lm(c(200,201,180,85,99,87)~c(1,1,1,0,0,0))) +  
  (mean(c(85,99,87)) + mean(c(200,201,180))) / 2
```

```
##           1           2           3           4           5           6  
## 148.3333 149.3333 128.3333 136.6667 150.6667 138.6667
```

# Correct with Combat

```
library("sva")  
ComBat(expr, batch)
```

```
## Standardizing Data across genes
```

```
##           S1           S2           S3           S4           S5           S6  
## G1 147.2891 147.9344 134.3816 134.2436 152.8668 136.9041  
## G2 113.6427 132.9256 139.3532 120.9387 139.7290 133.4656
```

# How does Combat work? Idea of Frequentist Fitting

Combat uses Bayesian fitting of linear model.

In contrast, Limma uses Frequentist fitting which is based on Maximum Likelihood principle:

$$y = \alpha + \beta x$$

$$L(y) \sim e^{-\frac{(y-\alpha-\beta x)^2}{2\sigma^2}}$$

$$\max_{\alpha, \beta, \sigma} L(y) \implies \hat{\alpha}, \hat{\beta}, \hat{\sigma}$$

## How does Combat work? Idea of Bayesian Fitting

Bayesian fitting is based on Maximum Posterior Probability principle:

$$y \sim N(\mu, \sigma) \quad \text{- Likelihood } L(y)$$

$$\mu = \alpha + \beta x$$

$$\alpha \sim N(\mu_\alpha, \sigma_\alpha) \quad \text{- Prior on } \alpha$$

$$\beta \sim N(\mu_\beta, \sigma_\beta) \quad \text{- Prior on } \beta$$

$$P(\mu_\alpha, \sigma_\alpha, \mu_\beta, \sigma_\beta, \sigma) \sim L(y) * N(\mu_\alpha, \sigma_\alpha) * N(\mu_\beta, \sigma_\beta)$$

$$\max_{\mu_\alpha, \sigma_\alpha, \mu_\beta, \sigma_\beta, \sigma} P(\mu_\alpha, \sigma_\alpha, \mu_\beta, \sigma_\beta, \sigma) \implies \hat{\mu}_\alpha, \hat{\sigma}_\alpha, \hat{\mu}_\beta, \hat{\sigma}_\beta, \hat{\sigma}$$

All genes “know” about each other, CIs are smaller than by one-by-one fitting - Bayesian shrinkage toward the mean (Random Effects modeling in Frequentist).

## How does Combat work? Bayesian Fitting

```
join<-data.frame(Expr=c(expr["G1",],expr["G2",]),Batch=c(batch, batch),  
                 Gene_ID=c(rep(1,6),rep(2,6)))
```

```
library("brms")
```

```
brmfit <- brm(Expr ~ Batch + (Batch | Gene_ID), data = join)
```

```
newvary <- expand.grid(Batch = 0:1, Gene_ID = 1:2)
```



```
par(mfrow=c(1,2))
```

```
plot(G1~batch,data=df); abline(lm(G1~batch,data=df))
```

```
lines(predict(brmfit, newvary)[1:2]~c(0,1),col="blue")
```


```
plot(G2~batch,data=df); abline(lm(G2~batch,data=df))
```

```
lines(predict(brmfit, newvary)[3:4]~c(0,1),col="red")
```



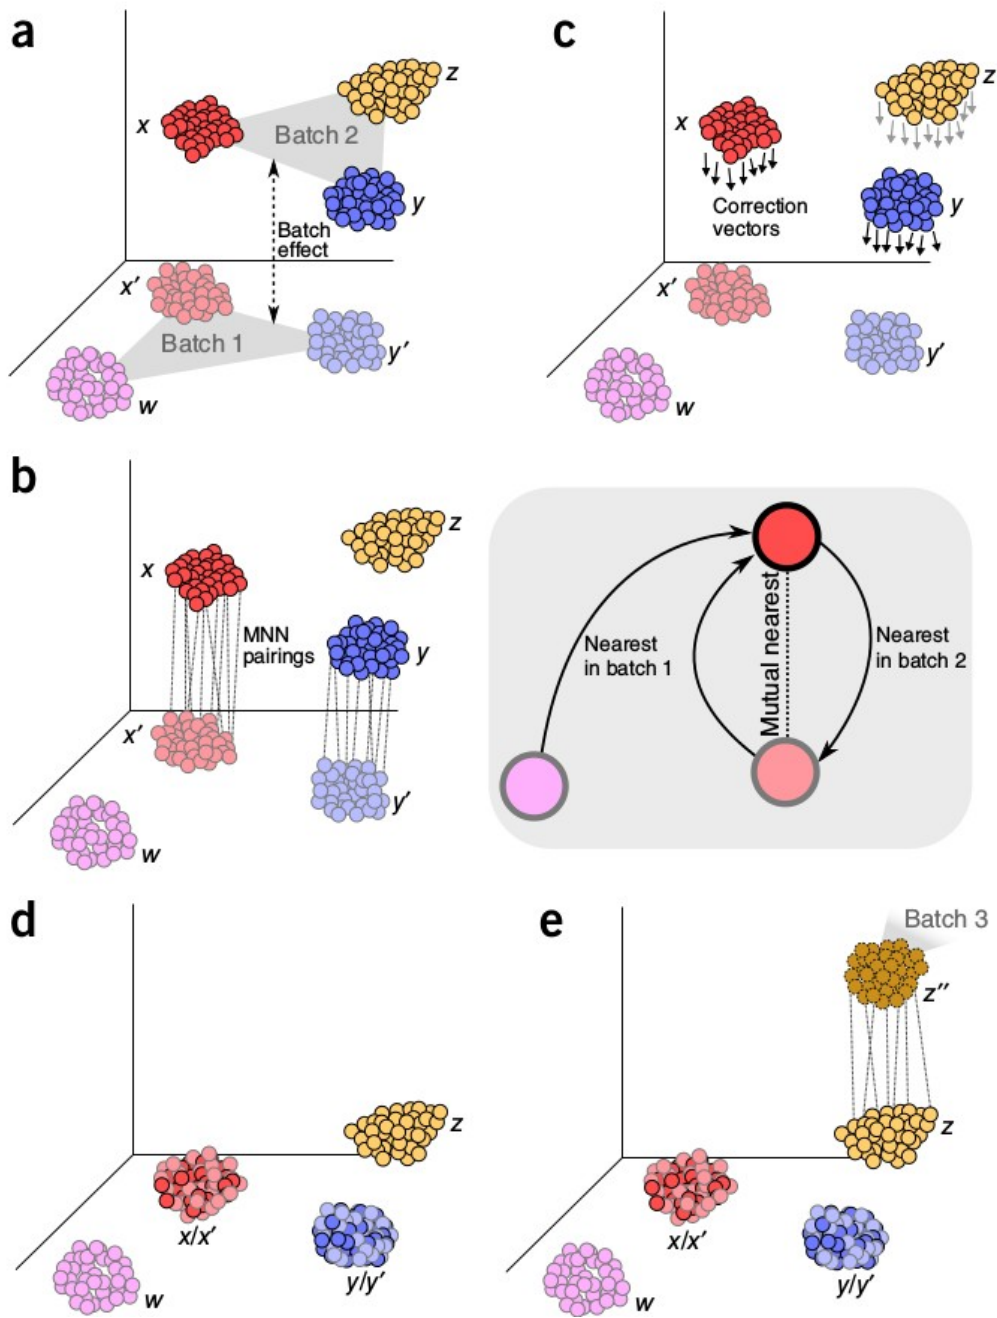
# Bayesian Variance Stabilization

```
newvary <- expand.grid(Gene_ID=1:2, Batch=seq(from=0, by=0.1, to=1))
fitvary <- cbind(newvary, fitted(brmfit, newdata = newvary)[, -2])
names(fitvary) <- c("Gene_ID", "Batch", "Expr", "lower", "upper")
p <- ggplot(join, aes(x=Batch, y=Expr)) + geom_point() +
  geom_smooth(method="lm")
p_sep <- p + facet_wrap(~Gene_ID, nrow = 1)
p_sep + geom_line(data = fitvary, aes(y = Expr), size = 1) +
  geom_line(data = fitvary, aes(y = lower), lty = 2) +
  geom_line(data = fitvary, aes(y = upper), lty = 2)
```

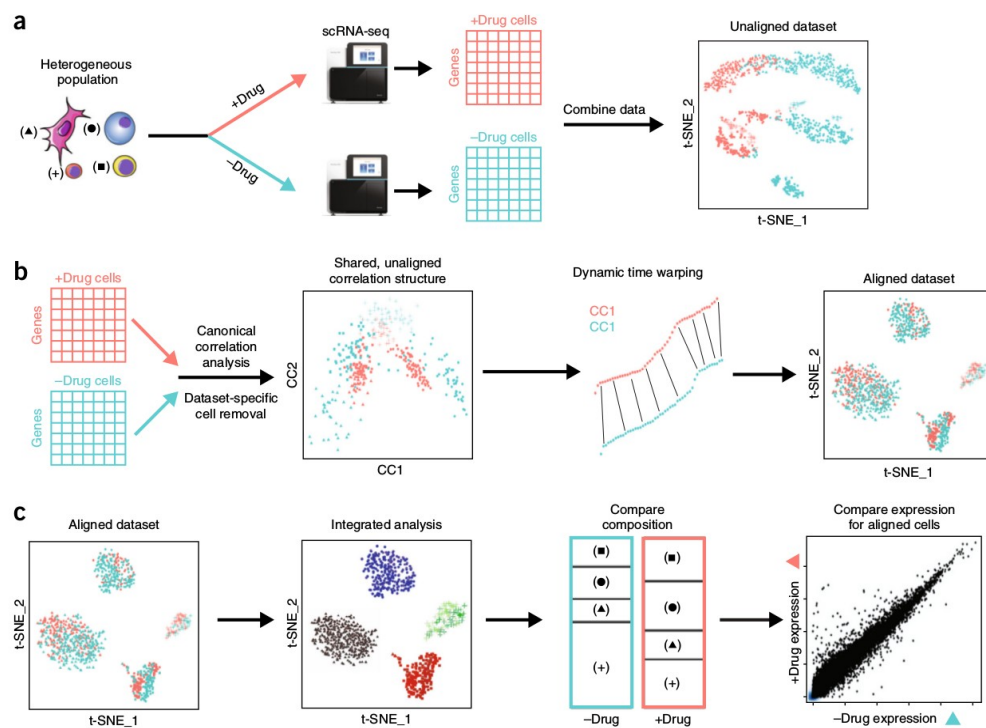


# Other Methods for Batch Effects Corrections


## MNN



## SEURAT



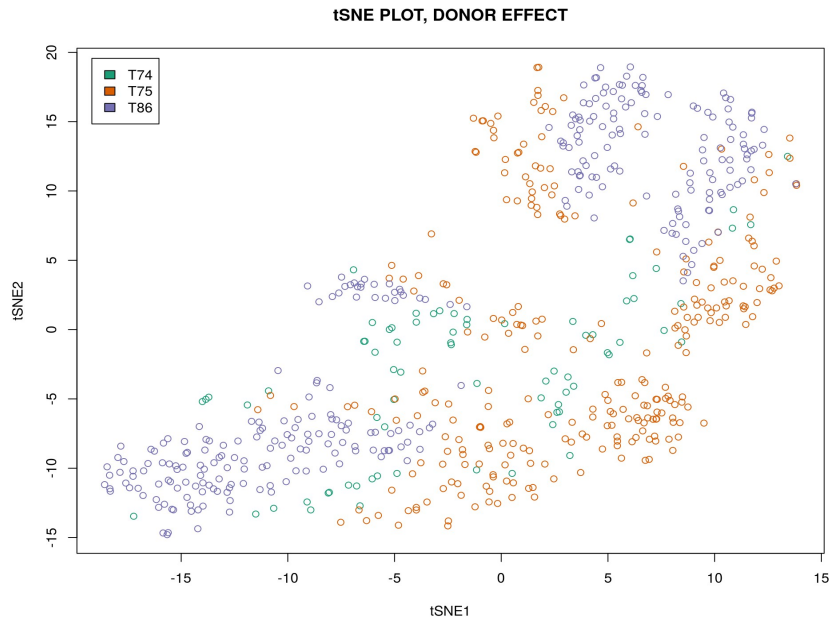
## SCMAP



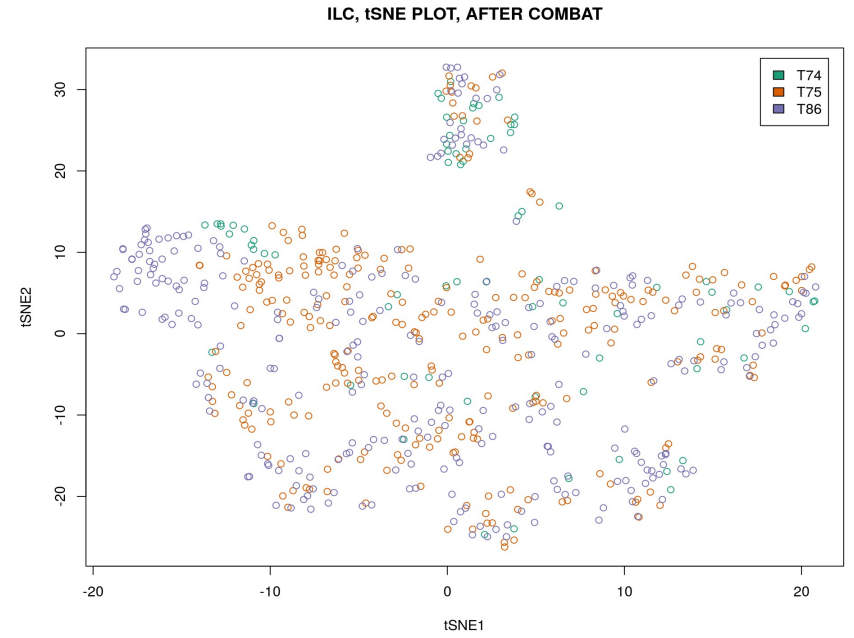


# Other Methods: Do They Work?

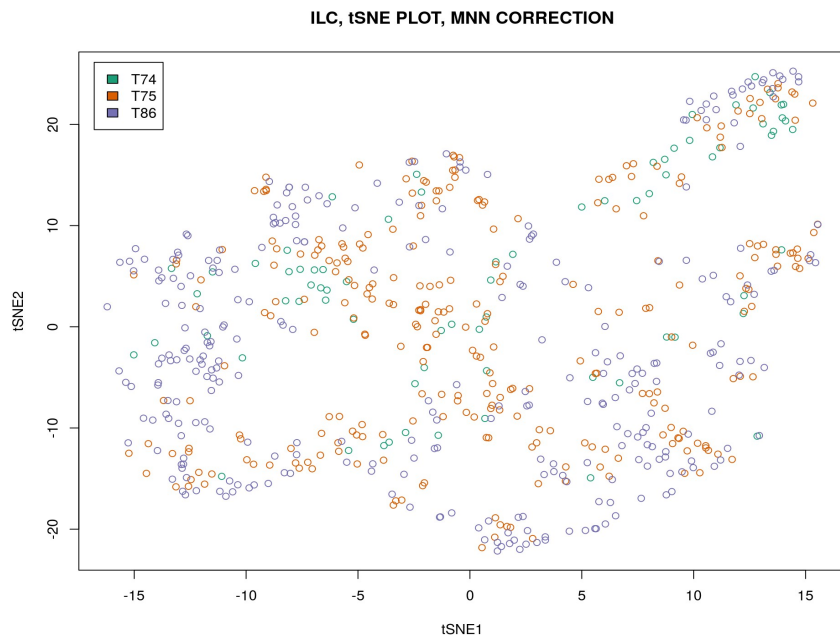
## UNCORRECTED




## COMBAT



## MNN



## SEURAT



# SCMAP

```
results_assigned<-results[as.character(results$ASSIGNED_LABEL)!="unassigned",]  
head(results_assigned,20)
```

```
##           CELS ASSIGNED_LABEL TRUE_LABEL SIMILARITY  
## 2  T86_P1_A10_ILC3           ILC3      ILC3  0.4457213  
## 3  T86_P1_A12_ILC3           ILC3      ILC3  0.4963317  
## 7   T86_P1_B1_ILC3           ILC3      ILC3  0.4955753  
## 8  T86_P1_B12_ILC3          ILC3      ILC3  0.4329544  
## 14   T86_P1_B9_NK            NK        NK  0.5676746  
## 18 T86_P1_C12_ILC3          ILC3      ILC3  0.5136711  
## 23  T86_P1_C6_ILC3          ILC3      ILC3  0.4655949  
## 28 T86_P1_D10_ILC3          ILC2      ILC3  0.3970456  
## 29   T86_P1_D11_NK          NK        NK  0.5262334  
## 30 T86_P1_D12_ILC3          ILC3      ILC3  0.5097175  
## 34  T86_P1_D6_ILC3          ILC3      ILC3  0.4664650  
## 38 T86_P1_E10_ILC3          ILC3      ILC3  0.4750463  
## 40 T86_P1_E12_ILC3          ILC3      ILC3  0.4968623  
## 41  T86_P1_E2_ILC3          ILC3      ILC3  0.4253116  
## 42  T86_P1_E3_ILC3          ILC3      ILC3  0.4706919  
## 45   T86_P1_E6_NK            NK        NK  0.5219235  
## 46   T86_P1_E7_NK            NK        NK  0.5405412  
## 48  T86_P1_E9_ILC3          ILC3      ILC3  0.4489822  
## 50 T86_P1_F10_ILC3          ILC3      ILC3  0.4821195  
## 51 T86_P1_F11_ILC3          ILC3      ILC3  0.4667251
```

```
table(results_assigned$ASSIGNED_LABEL,results_assigned$TRUE_LABEL)
```

```
##  
##           ILC1 ILC2 ILC3 NK  
## ILC1           65  26  0  0  
## ILC2            0   2  2  0  
## ILC3            0   0 60  0  
## NK              0   0  0 18  
## unassigned     0   0  0  0
```

```
sum(as.character(results_assigned$ASSIGNED_LABEL)==as.character(results_assigned$TRUE_LABEL))/dim(results_assigned)[1]
```

```
## [1] 0.8381503
```

We conclude that the accuracy of assignment is 84% which is not fantastic taking into account that SCMAP failed assignment of almost a half of the cells in the test data set.

# **Brief Overview of Bulk RNAseq Normalization Methods:**

**RPKM, DESeq / TMM**

# RPKM<sub>s</sub> (FPKM<sub>s</sub>)

RPKM normalization is an extension of so-called library size normalization

Library size normalization: scaling such that library size is equal between all libraries

$$RPKM = \frac{10^9 C}{NL}$$

where:

$C$  = number of reads that overlap a given gene

$N$  = library size

$L$  = gene length

Disadvantage: forced equalizing library sizes might eliminate true biological variation

# SimplexSpace

Nikolay Oskolkov

May 5, 2018

## Generate Gene Expression

We draw two vectors (genes) from normal distributions with different means and sds

```
gene1<-rnorm(100,mean=10,sd=1)  
gene2<-rnorm(100,mean=100,sd=5)
```

## The Genes are Uncorrelated (Pearson)

We make sure that the two genes are not correlated

```
cor.test(gene1, gene2, method="pearson")
```

```
##
```

```
## Pearson's product-moment correlation
```

```
##
```

```
## data: gene1 and gene2
```

```
## t = -0.0058837, df = 98, p-value = 0.9953
```

```
## alternative hypothesis: true correlation is not equal to 0
```

```
## 95 percent confidence interval:
```

```
## -0.1969895 0.1958466
```

```
## sample estimates:
```

```
##          cor
```

```
## -0.000594346
```

Pearson correlation coefficient is close to 0

## Library Size Normalization

We create an expression data frame with 100 samples and 2 genes

```
expr<-as.data.frame(t(data.frame(gene1=gene1,gene2=gene2)))  
colnames(expr)<-paste0("sample",seq(1:100))  
expr[1:2,1:4]
```

```
##           sample1    sample2    sample3    sample4  
## gene1  8.969518  11.07167    8.189429    9.838804  
## gene2 96.493377 107.67880 102.187541 107.134326
```

We calculate library size normalized expression data frame

```
expr_sim<-apply(expr,2,function(x) x/sum(x))  
expr_sim[1:2,1:4]
```

```
##           sample1    sample2    sample3    sample4  
## gene1 0.08504904 0.09323475 0.07419508 0.08411165  
## gene2 0.91495096 0.90676525 0.92580492 0.91588835
```



## Spurious Pearson Correlation After Normalization

We prove that before library size normalization, the genes were uncorrelated but after the normalization they become 100% negatively correlated

```
cor.test(as.numeric(expr_sim["gene1",]),  
         as.numeric(expr_sim["gene2",]),method="pearson")
```

```
##
```

```
## Pearson's product-moment correlation
```

```
##
```

```
## data: as.numeric(expr_sim["gene1", ]) and as.numeric(expr_sim["gene
```

```
## t = -664340000, df = 98, p-value < 2.2e-16
```

```
## alternative hypothesis: true correlation is not equal to 0
```

```
## 95 percent confidence interval:
```

```
## -1 -1
```


```
## sample estimates:
```

```
## cor
```

```
## -1
```


# Plot Genes Before Normalization

```
plot(as.numeric(expr["gene1",]),as.numeric(expr["gene2",]))
```



# Plot Genes After Normalization

```
plot(as.numeric(expr_sim["gene1",]),as.numeric(expr_sim["gene2",]))
```



# DESeq

DESeq: create reference library based on geometric mean of all libraries, calculate size factors as ratios against the reference library

$$\begin{array}{c}
 \mathbf{g}_1 \\
 \mathbf{g}_2 \\
 \mathbf{g}_n
 \end{array}
 \begin{array}{ccc}
 \mathbf{s}_1 & \mathbf{s}_2 & \mathbf{s}_k \\
 \left[ \begin{array}{ccc}
 c_{11} & c_{12} & c_{1k} \\
 c_{21} & c_{22} & c_{2k} \\
 c_{n1} & c_{n2} & c_{nk}
 \end{array} \right]
 \end{array}
 \begin{array}{c}
 \left[ \begin{array}{ccc}
 1/g_{1r} & 1/g_{2r} & 1/g_{nr}
 \end{array} \right]
 \end{array}
 \longrightarrow
 \begin{array}{c}
 \mathbf{g}_1 \\
 \mathbf{g}_2 \\
 \mathbf{g}_n
 \end{array}
 \begin{array}{ccc}
 \mathbf{s}_1 & \mathbf{s}_2 & \mathbf{s}_k \\
 \left[ \begin{array}{ccc}
 r_{11} & r_{12} & r_{1k} \\
 r_{21} & r_{22} & r_{2k} \\
 r_{n1} & r_{n2} & r_{nk}
 \end{array} \right]
 \end{array}
 \begin{array}{l}
 \text{estimate the relative} \\
 \text{depth of the library} \\
 \\
 S_1 = \text{median}(r_{x1})
 \end{array}$$

take the mean for each row to obtain a reference sample

$$\begin{bmatrix} g_{1r} \\ g_{2r} \\ g_{nr} \end{bmatrix}$$

estimate the depth ratio for each gene

$$\begin{array}{c}
 \mathbf{g}_1 \\
 \mathbf{g}_2 \\
 \mathbf{g}_n
 \end{array}
 \begin{array}{ccc}
 \mathbf{s}_1 & \mathbf{s}_2 & \mathbf{s}_k \\
 \left[ \begin{array}{ccc}
 c_{11} & c_{12} & c_{1k} \\
 c_{21} & c_{22} & c_{2k} \\
 c_{n1} & c_{n2} & c_{nk}
 \end{array} \right]
 \end{array}
 \begin{array}{c}
 \left[ \begin{array}{c}
 S_1 \\
 S_2 \\
 S_k
 \end{array} \right]
 \end{array}$$



Disadvantage: DESeq is based on ratio construction

**scRNAseq – Specific Normalization Methods:  
Deconvolution (Pooling-Across-Cells-Method)**

# Lots of zero-counts is main challenge in scRNAseq

scRNAseq expression counts have typically ~80% of zero-counts

This is due to: 1) low amounts of RNA per cell, 2) RNA capture efficiency



We want to correct for sequencing depth and cell-to-cell difference in RNA capture efficiency

3 common normalization methods used for bulk RNAseq: 1) TMM, 2) DESeq, 3) RPKM

Main assumption of all 3 methods: most of the genes are not differentially expressed

TMM and DESeq rely on ratios of counts, therefore diverge when lots of zero-counts

# Deconvolution Normalization Method

METHOD

Open Access



## Pooling across cells to normalize single-cell RNA sequencing data with many zero counts

Aaron T. L. Lun<sup>1\*</sup>, Karsten Bach<sup>2</sup> and John C. Marioni<sup>1,2,3\*</sup>

**Abstract**

Normalization of single-cell RNA sequencing data is necessary to eliminate cell-specific biases prior to downstream analyses. However, this is not straightforward for noisy single-cell data where many counts are zero. We present a novel approach where expression values are summed across pools of cells, and the summed values are used for normalization. Pool-based size factors are then deconvolved to yield cell-based factors. Our deconvolution approach outperforms existing methods for accurate normalization of cell-specific biases in simulated data. Similar behavior is observed in real data, where deconvolution improves the relevance of results of downstream analyses.

**Keywords:** Single-cell RNA-seq, Normalization, Differential expression

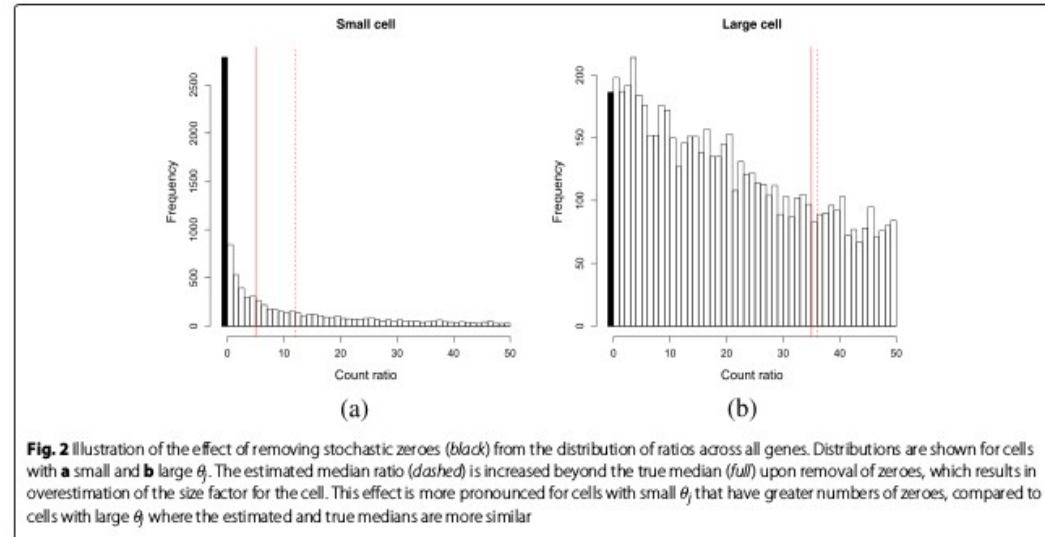
**Background**

Single-cell RNA sequencing (scRNA-seq) is a powerful technique that allows researchers to characterize the gene expression profile of single cells. From each cell, mRNA is isolated and reverse-transcribed into cDNA, which is amplified and subjected to massively parallel sequencing [1]. The sequencing reads are mapped to a reference genome, such that the number of reads mapped to each gene can be used to quantify its expression. Alternatively, transcript molecules can be counted directly using unique molecular identifiers (UMIs) [2]. Count data can be analyzed to identify new cell subtypes and to detect highly variable or differentially expressed (DE) genes between cell subpopulations. This type of single-cell resolution is not possible with bulk RNA sequencing of cellular populations. However, the downside is that the counts often contain high levels of technical noise with many dropouts, i.e., zero or near-zero values. This is due to the presence of low amounts of RNA per cell, which decreases the efficiency with which transcripts can be captured and processed prior to sequencing. Moreover, the capture

efficiency often varies from cell to cell, such that counts cannot be directly compared between cells.

Normalization of the scRNA-seq counts is a critical step that corrects for cell-to-cell differences in capture efficiency, sequencing depth, and other technical confounders. This ensures that downstream comparisons of relative expression between cells are valid. Two broad classes of methods for scaling normalization are available: those using spike-in RNA sets and those using the counts from the profiled cellular RNA. In the former, the same quantity of spike-in RNA is added to each cell prior to library preparation [1]. Any difference in the coverage of the spike-in transcripts must be caused by differences in capture efficiency, amplification bias, or sequencing depth between cells. Normalization is then performed by scaling the counts to equalize spike-in coverage between cells. For the methods using cellular counts, the assumption is that most genes are not DE across the sampled cells. Counts are scaled so that there is, on average, no fold-difference in expression between cells for the majority of genes. This is the underlying concept of commonly used methods such as DESeq [3] and trimmed mean of *M* values (TMM) normalization [4]. An even simpler approach involves scaling the counts to remove differences in library sizes between cells, i.e., library size normalization.

The type of normalization that can be used depends on the characteristics of the data set. In some cases, spike-in



of an arbitrary set of cells  $S_k$ . Define  $V_{ik}$  as the sum of  $Z_{ij}$  across all cells in  $S_k$ , which has an expectation of

$$E(V_{ik}) = \lambda_{i0} \sum_{j \in S_k} \theta_j t_j^{-1}.$$


The observed values of  $V_{ik}$  across all genes constitute an overall expression profile for the pool of cells corresponding to  $S_k$ . Also define  $U_i$  as the mean of  $Z_{ij}$  across all  $N$  cells in the entire data set, which has an expectation of

$$E(U_i) = \lambda_{i0} N^{-1} \sum_{j \in S_0} \theta_j t_j^{-1}$$

where  $S_0$  refers to the set of all cells in the data set. The observed values of  $U_i$  across all genes represent the expression profile for an averaged reference pseudo-cell.

The cell pool  $k$  is then normalized against this reference pseudo-cell. Define  $R_{ik}$  as the ratio of  $V_{ik}$  to  $U_i$  for the non-DE gene  $i$ . The expectation of  $R_{ik}$  represents the true size factor for the pooled cells in  $S_k$ , and is written as

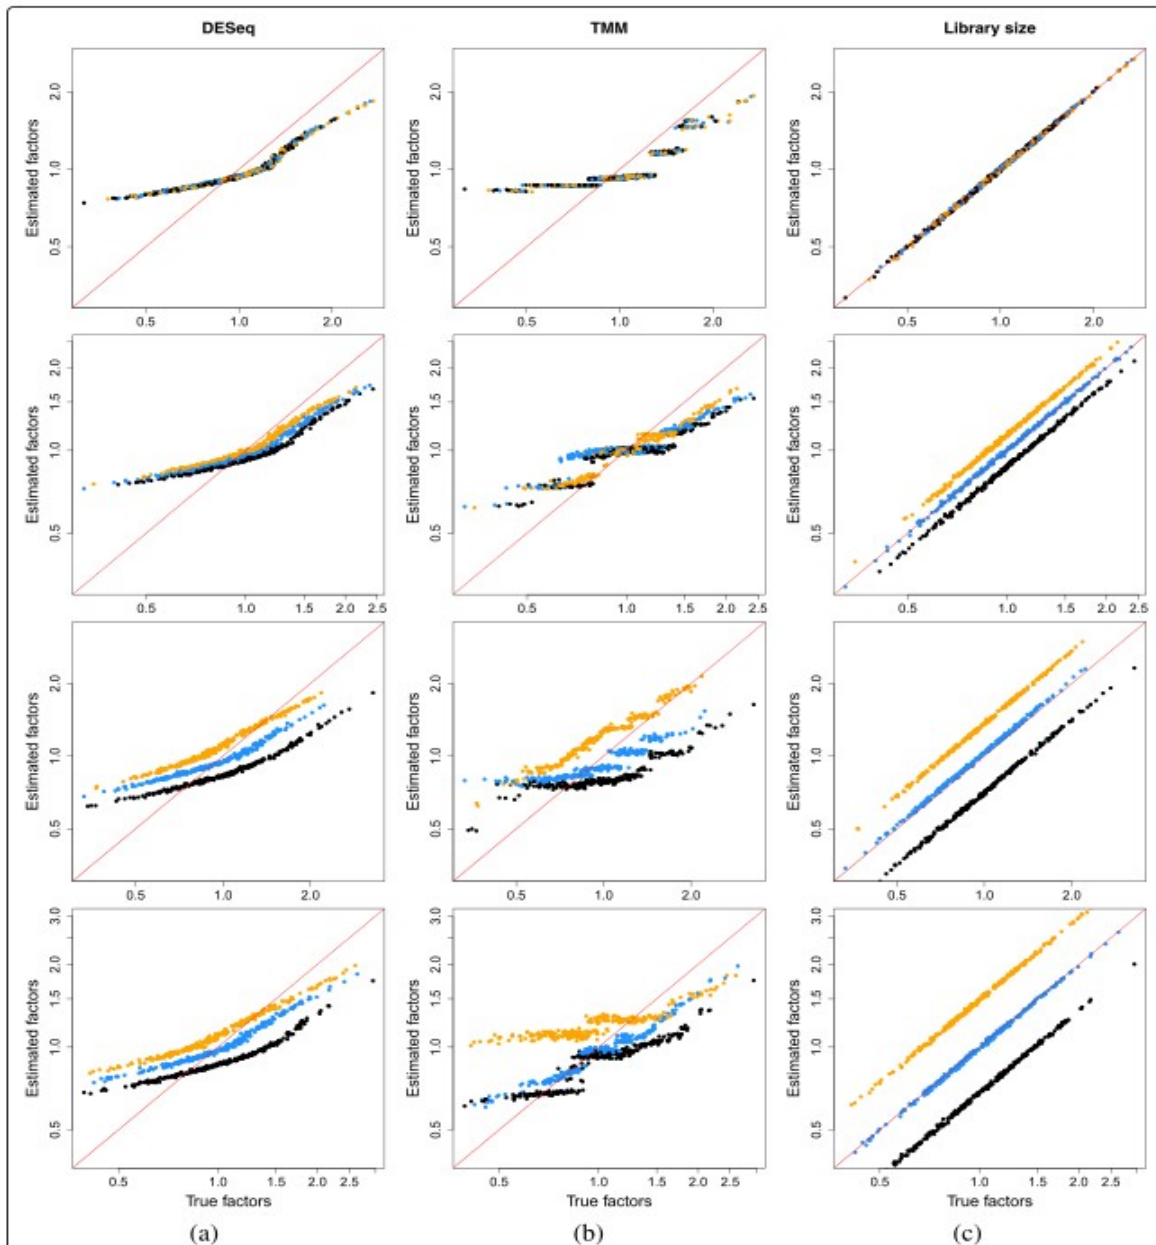
$$E(R_{ik}) \approx \frac{E(V_{ik})}{E(U_i)} = \frac{\sum_{j \in S_k} \theta_j t_j^{-1}}{N^{-1} \sum_{j \in S_0} \theta_j t_j^{-1}} = \frac{\sum_{j \in S_k} \theta_j t_j^{-1}}{C} \quad (1)$$



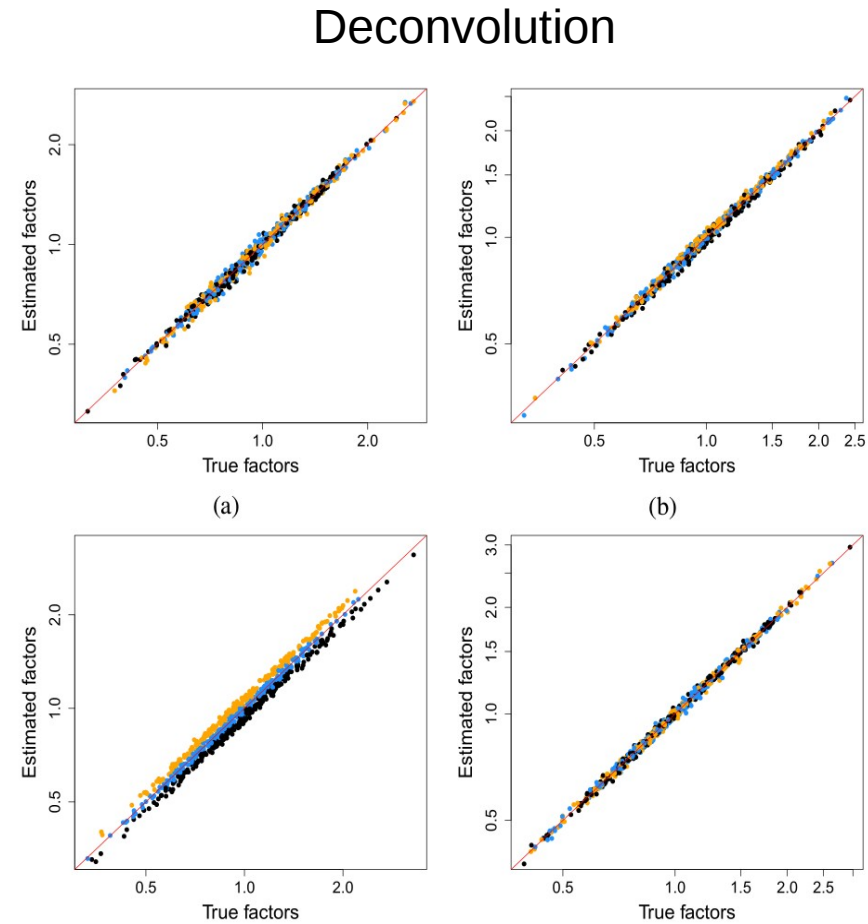
**Fig. 3** Schematic of the deconvolution method. All cells in the data set are averaged to make a reference pseudo-cell. Expression values for cells in pool A are summed together and normalized against the reference to yield a pool-based size factor  $\theta_A$ . This is equal to the sum of the cell-based factors  $\theta_j$  for cells  $j = 1-4$  and can be used to formulate a linear equation. (For simplicity, the  $t_j$  term is assumed to be unity here.) Repeating this for multiple pools (e.g., pool B) leads to the construction of a linear system that can be solved to estimate  $\theta_j$  for each cell  $j$

\*Correspondence: aaron.lun@cruk.cam.ac.uk; marioni@ebi.ac.uk  
<sup>1</sup> Cancer Research UK Cambridge Institute, University of Cambridge, Li Ka Shing Centre, Robinson Way, CB2 0RE, Cambridge, UK  
<sup>2</sup> EMBL European Bioinformatics Institute, Wellcome Genome Campus, Hinxton, CB10 1SD, Cambridge, UK  
 Full list of author information is available at the end of the article

# Benchmarking: Deconvolution Method Performs Best





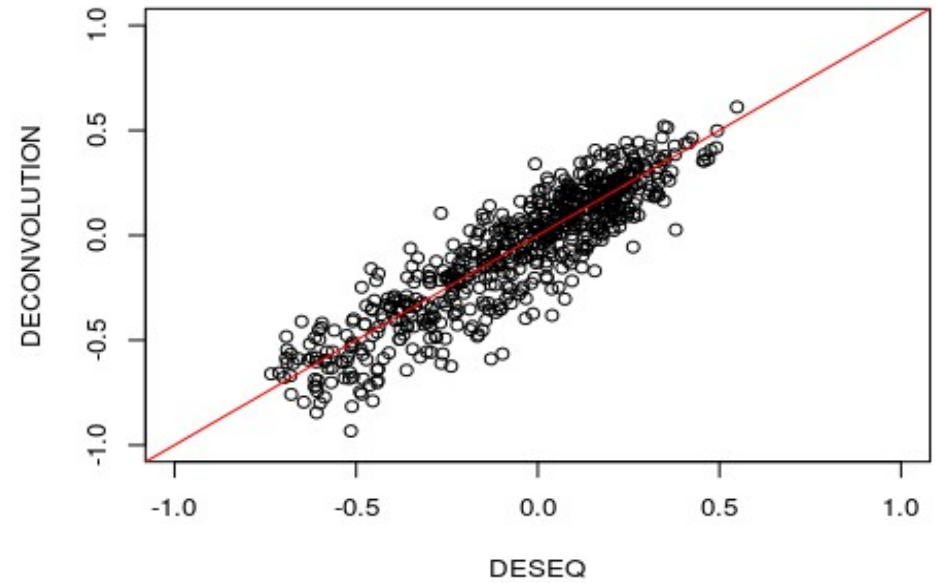

**Fig. 1** Performance of existing normalization methods on the simulated data with DE genes and stochastic zeroes. The size factor estimates for all cells are plotted against the true values for **a** DESeq, **b** TMM, and **c** library size normalization. Simulations were performed with no DE (*first row*), moderate DE (*second row*), strong DE (*third row*), and varying magnitudes of DE (*fourth row*). Axes are shown on a log-scale. For comparison, each set of size factors was scaled such that the grand mean across cells was the same as that for the true values. The red line represents equality between the rescaled estimates and true factors. Cells in the first, second, and third subpopulations are shown in black, blue, and orange, respectively. DE differentially expressed, TMM trimmed mean of  $M$  values





# Deconvolution vs TMM vs DESeq vs RPKM: Size Factors

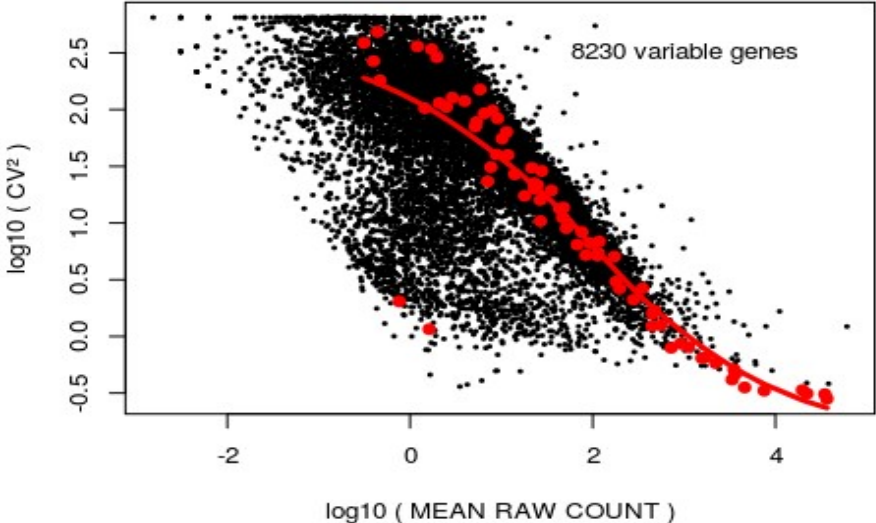
For other data sets it might not look as good as for ILC!



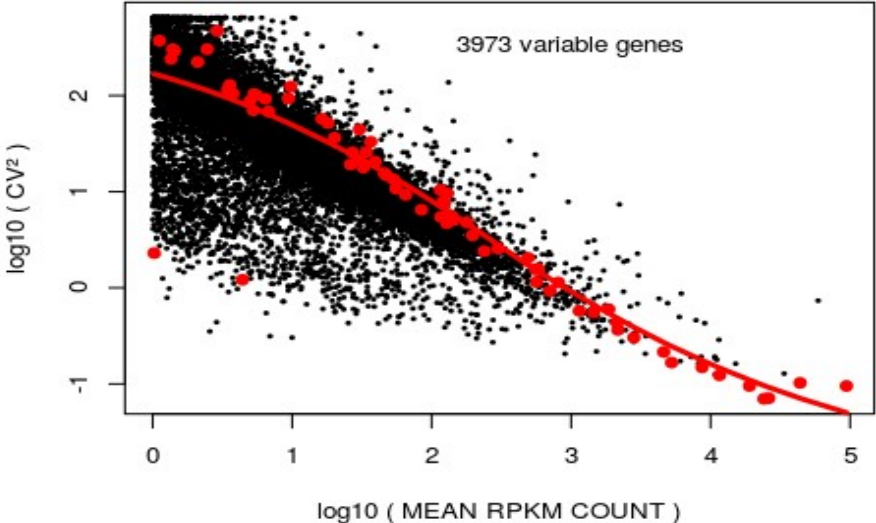
**How does deconvolution normalization method  
compare with RPKM and normalizations  
by using spike-ins?**

# CV<sup>2</sup> vs. Mean Expression Plot


### RAW COUNTS




### RPKM COUNTS



### DECONVOLUTION COUNTS




### SPIKE LIB SIZE COUNTS




# PCA Plot


### PCA PLOT: RAW COUNTS




### PCA PLOT: RPKM COUNTS



### PCA PLOT: DECONVOLUTION COUNTS




### PCA PLOT: SPIKE LIB SIZE COUNTS




# tSNE Plot


### tSNE: RAW COUNTS




### tSNE: RPKM COUNTS



### tSNE: DECONVOLUTION COUNTS




### tSNE: SPIKE LIB SIZE COUNTS




# Cell Cycle Phase Assignment

Pre-trained classifier looks at pairs of genes having difference in expression that changes sign from phase to phase of cell cycle

PCA PLOT: DECONVOLUTION COUNTS



tSNE: DECONVOLUTION COUNTS



# Methods for Testing for Differential Expression without Normalization:

**SCDE**

## Bayesian approach to single-cell differential expression analysis

Peter V Kharchenko<sup>1-3</sup>, Lev Silberstein<sup>3-5</sup> & David T Scadden<sup>3-5</sup>

Single-cell data provide a means to dissect the composition of complex tissues and specialized cellular environments. However, the analysis of such measurements is complicated by high levels of technical noise and intrinsic biological variability. We describe a probabilistic model of expression-magnitude distortions typical of single-cell RNA-sequencing measurements, which enables detection of differential expression signatures and identification of subpopulations of cells in a way that is more tolerant of noise.

Methodological advances are making it possible to examine transcription in individual cells on a large scale<sup>1-4</sup>, facilitating unbiased analysis of cellular states<sup>5-8</sup>. However, profiling the low amounts of mRNA within individual cells typically requires amplification by more than 1 million fold, which leads to severe nonlinear distortions of relative transcript abundance and accumulation of nonspecific byproducts. A low starting amount also makes it more likely that a transcript will be 'missed' during the reverse-transcription step and consequently not detected during sequencing. This leads to so-called 'dropout' events, in which a gene is observed at a moderate or high expression level in one cell but is not detected in another cell (Fig. 1a). More fundamentally, gene expression is inherently stochastic, and some cell-to-cell variability will be an unavoidable consequence of transcriptional bursts of individual genes or coordinated fluctuations of multigene networks<sup>9</sup>. Such biological variability is of high interest, and several methods have been proposed for detecting it<sup>10-12</sup>. Collectively, this multifactorial variability in single-cell measurements substantially increases the apparent level of noise, posing challenges for differential expression and other downstream analyses.

Comparisons of RNA-seq data from individual cells tend to show higher variability than is typically observed in biological replicates of bulk RNA-seq measurements. In addition to strong overdispersion, there are high-magnitude outliers as well as dropout events (Fig. 1a). Such variability is poorly accommodated by

standard RNA-seq analysis methods<sup>13,14</sup>, and the reported sets of top differentially expressed genes can include high-magnitude outliers or dropout events, showing poor consistency within each cell population (Fig. 1b). The abundance of dropout events has been previously noted in single-cell quantitative PCR data and accommodated with zero-inflated distributions<sup>15</sup>.

Two prominent characteristics of dropout events make them informative in further analysis of expression state. First, the overall dropout rates are consistently higher in some single-cell samples than in others (Supplementary Figs. 1 and 2), indicating that the contribution of an individual sample to the downstream cumulative analysis should be weighted accordingly. Second, the dropout rate for a given cell depends on the average expression magnitude of a gene in a population, with dropouts being more frequent for genes with lower expression magnitude. Quantification of such dependency provides evidence about the true expression magnitude. For instance, dropout of a gene observed at very high expression magnitude in other cells is more likely to be indicative of true expression differences than of stochastic variability.

We modeled the measurement of each cell as a mixture of two probabilistic processes—one in which the transcript is amplified and detected at a level correlating with its abundance and the other in which a transcript fails to amplify or is not detected for other reasons. We modeled the first, 'correlated' component with a negative binomial distribution<sup>13,16</sup>. The RNA-seq signal associated with the second, dropout component could in principle be modeled as a constant zero (i.e., zero-inflated negative binomial process); however, we used a low-magnitude Poisson process to account for some background signal that is typically detected for the dropout and transcriptionally silent genes. Importantly, the mixing ratio between the correlated and dropout processes depends on the magnitude of gene expression in a given cell population. We analyzed two single-cell data sets—a 92-cell set consisting of mouse embryonic fibroblast (MEF) and embryonic stem (ES) cells<sup>2</sup> and a data set of cells from different stages of early mouse embryos<sup>12</sup>. To fit the parameters of an error model for a particular single-cell measurement, we used a subset of genes for which an expected expression magnitude within the cell population can be reliably estimated. Briefly, we analyzed pairs of all other single-cell samples from the same subpopulation (for example, all MEF cells except for the one being fit) with a similarly structured three-component mixture containing one correlated component and dropout components for each cell (Fig. 1c and Supplementary Figs. 1 and 2). We deemed a subset of genes appearing in correlated components in a sufficiently large fraction of pairwise cell comparisons to be reliable. We estimated the expected expression magnitude of these

<sup>1</sup>Center for Biomedical Informatics, Harvard Medical School, Boston, Massachusetts, USA. <sup>2</sup>Hematology/Oncology Program, Children's Hospital, Boston, Massachusetts, USA. <sup>3</sup>Harvard Stem Cell Institute, Cambridge, Massachusetts, USA. <sup>4</sup>Center for Regenerative Medicine, Massachusetts General Hospital, Boston, Massachusetts, USA. <sup>5</sup>Department of Stem Cell and Regenerative Biology, Harvard University, Cambridge, Massachusetts, USA. Correspondence should be addressed to P.V.K. (peter.kharchenko@post.harvard.edu).

RECEIVED 2 OCTOBER 2013; ACCEPTED 28 MARCH 2014; PUBLISHED ONLINE 18 MAY 2014; DOI:10.1038/NMETH.2967





# Single-Cell Differential Expression (SCDE) Method

