

CRISPR Screen and Gene Expression Differential Analysis

Lianbo Yu, Yue Zhao, Kevin R. Coombes, and Lang Li

2022-01-06

Contents

| | | |
|----------|-------------------------|----------|
| 1 | Introduction | 1 |
| 2 | Overview | 1 |
| 2.1 | Data Format | 1 |
| 2.2 | Normalization | 2 |
| 2.3 | Analysis | 3 |

1 Introduction

We developed CEDA to analyze read counts of single guide RNAs (sgRNAs) raw CRISPR screening experiments. The sgRNAs are synthetically generated from genes, and each gene can generate multiple sgRNAs. CEDA models the sgRNA counts at different levels of gene expression by multi-component normal mixtures, with the model fit by an EM algorithm. Posterior estimates at sgRNA level are then summarized for each gene.

In this document, we use data from an experiment with the MDA231 cell line to illustrate how to use CEDA to perform CRISPR screen data analysis.

2 Overview

CEDA analysis follows a workflow that is typical for most omics level experiments.

1. Put the data into an appropriate format for input to CEDA.
2. Normalize the raw counts.
3. Fit a linear model to the data.
4. Summarize and view the results.

2.1 Data Format

In our experiment, three samples of MDA231 cells were untreated at time $T=0$, and another three samples of MDA231 cells were treated with DMSO at time $T=0$. We are interested in detecting sgRNAs that are differentially changed by a treatment.

The sgRNA read counts, along with a list of non-essential genes, are stored in the dataset `mda231` that we have included in the CEDA package. We read that dataset and explore its structure.

```
library(CEDA)
data("mda231")
class(mda231)
#> [1] "list"
length(mda231)
```

```
#> [1] 2
names(mda231)
#> [1] "sgRNA" "neGene"
```

As you can see, this is a list containing two components

1. `sgRNA`, the observed count data of six samples, and
2. `neGene`, the set of non-essential genes.

```
dim(mda231$sgRNA)
#> [1] 23618 9
length(mda231$neGene$Gene)
#> [1] 200
head(mda231$sgRNA)
#>
#>      sgRNA  Gene DMSOa DMSOb DMSOc TOa TOb TOc
#> 64780 chr19:10655652-10655671_ATG4D_ - ATG4D 126 100 132 94 82 78
#> 67381 chr5:32739109-32739128_NPR3_ + NPR3 266 452 309 557 687 587
#> 67411 chr3:45515731-45515750_LARS2_ + LARS2 28 45 36 583 660 512
#> 27053 chr12:111856039-111856058_SH2B3_ + SH2B3 509 501 661 578 824 636
#> 55806 chr9:118163517-118163536_DEC1_ + DEC1 265 489 390 718 733 655
#> 57274 chr1:228879268-228879287_RHOV_ + RHOV 144 124 137 160 164 119
#>      exp.level.log2
#> 64780 3.655866985
#> 67381 0.200236907
#> 67411 3.495375381
#> 27053 4.227348316
#> 55806 0.005794761
#> 57274 0.925993344
```

Notice that the `sgRNA` component includes an extra column, “`exp.level.log2`”, that are the expression level (in log2 scale) of genes and was computed from raw gene expression data.

The second element of the list `neGene` is, as expected, just a list of gene names that are the non-essential genes:

```
dim(mda231$neGene)
#> [1] 200 1
head(mda231$neGene)
#>      Gene
#> 189 MMD2
#> 303 SUN5
#> 155 KRT2
#> 72 DEFA5
#> 195 MUC17
#> 70 CYP2C19
```

2.2 Normalization

The `sgRNA` read counts needs to be normalized across sample replicates before formal analysis. The non-essential genes are assumed to have no change after DMSO treatment. So, our recommended procedure is to perform median normalization based on the set of non-essential genes. `nes` were used for normalizing `sgRNA` counts of all samples.

```
mda231.ne <- mda231$sgRNA[mda231$sgRNA$Gene %in% mda231$neGene$Gene,]
cols <- c(3:8)
mda231.norm <- medianNormalization(mda231$sgRNA[,cols], mda231.ne[,cols])[[2]]
```

2.3 Analysis

Our primary goal is to detect essential sgRNAs that have different count levels between conditions. We rely on the R package `limma` to calculate log fold ratios between three untreated and three treated samples.

2.3.1 Calculating fold ratios

First, we have to go through the usual `limma` steps to describe the design of the study. There were two groups of replicate samples. We will call these groups “Control” and “Baseline” (although “Treated” and “Untreated” would work just as well). Our main interest is determining the differences between the groups. And we have to record this information in a “contrast matrix” so `limma` knows what we want to compare.

```
group <- gl(2, 3, labels=c("Control", "Baseline"))
design <- model.matrix(~ 0 + group)
colnames(design) <- sapply(colnames(design), function(x) substr(x, 6, nchar(x)))
contrast.matrix <- makeContrasts("Control-Baseline", levels=design)
```

Finally, we can run the `limma` algorithm.

```
limma.fit <- runLimma(log2(mda231.norm+1), design, contrast.matrix)
```

We merge the results from our `limma` analysis with the original sgRNA count data.

```
mda231.limma <- data.frame(mda231$sgRNA, limma.fit)
head(mda231.limma)
#>
#>      sgRNA Gene DMSOa DMSOb DMSOc TOa TOb TOc
#> 64780 chr19:10655652-10655671_ATG4D_ - ATG4D 126 100 132 94 82 78
#> 67381 chr5:32739109-32739128_NPR3_ + NPR3 266 452 309 557 687 587
#> 67411 chr3:45515731-45515750_LARS2_ + LARS2 28 45 36 583 660 512
#> 27053 chr12:111856039-111856058_SH2B3_ + SH2B3 509 501 661 578 824 636
#> 55806 chr9:118163517-118163536_DEC1_ + DEC1 265 489 390 718 733 655
#> 57274 chr1:228879268-228879287_RHOV_ + RHOV 144 124 137 160 164 119
#>      exp.level.log2 lfc se p
#> 64780 3.655866985 0.52954908 0.1093812 6.287987e-03
#> 67381 0.200236907 -0.81763714 0.3240186 2.269143e-02
#> 67411 3.495375381 -3.94431072 0.2746740 1.624301e-06
#> 27053 4.227348316 -0.23583315 0.1257939 1.437600e-01
#> 55806 0.005794761 -0.87573339 0.3301200 1.830149e-02
#> 57274 0.925993344 -0.07125132 0.1276733 6.343237e-01
```

2.3.2 Fold ratios under the null hypotheses

Under the null hypotheses, all sgRNAs levels are unchanged between the two conditions. To obtain fold ratios under the null, samples were permuted between two conditions, and log fold ratios were obtained from `limma` analysis under each permutation.

```
betanull <- permuteLimma(log2(mda231.norm + 1), design, contrast.matrix, 20)
theta0 <- sd(betanull)
theta0
#> [1] 0.4500416
```

2.3.3 Fitting three-component mixture models

A three-component mixture model (unchanged, overexpressed, and underexpressed) is assumed for log fold ratios at different level of gene expression. Empirical Bayes method was employed to estimate parameters of the mixtures and posterior means were obtained for estimating actual log fold ratios between the two conditions. P-values of sgRNAs were then calculated by permutation method.

```
nmm.fit <- normalMM(mda231.limma, theta0)
```

Results from the mixture model were shown in Figure 1. False discovery rate of 0.05 was used for declaring significant changes in red color between the two conditions for sgRNAs.

```
scatterPlot(nmm.fit$data, fdr=0.05, xlim(-0.5, 12), ylim(-8, 5))
```

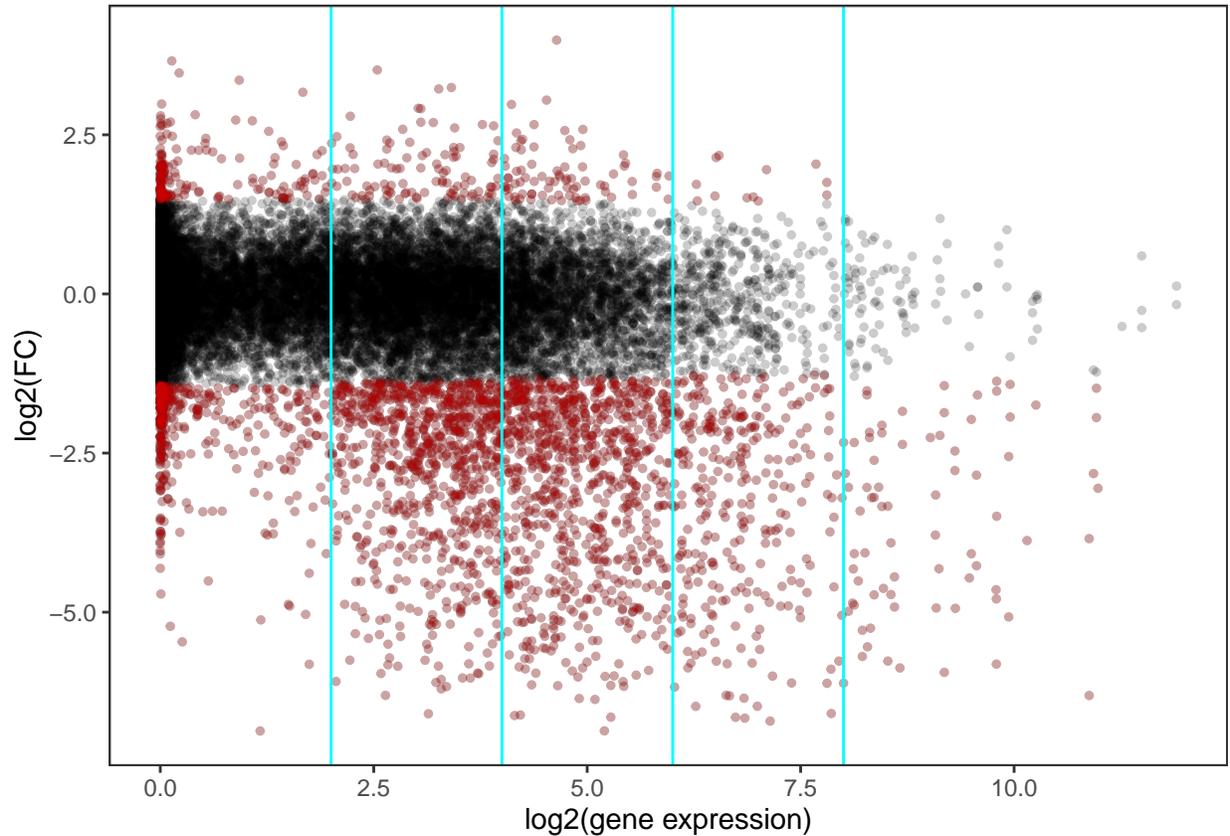


Figure 1: Log fold ratios of sgRNAs vs. gene expression level

2.3.4 Gene level summarization

From the p-values of sgRNAs, gene level p-values were obtained by using modified robust rank aggregation method (alpha-RRA). Log fold ratios were also summarized at gene level.

```
mda231.nmm <- nmm.fit[[1]]  
p.gene <- calculateGenePval(exp(mda231.nmm$log_p), mda231.nmm$Gene, 0.05)  
fdr.gene <- stats::p.adjust(p.gene$pvalue, method = "fdr")  
lfc.gene <- calculateGeneLFC(mda231.nmm$lfc, mda231.nmm$Gene)
```