

SIBER User Manual

Pan Tong and Kevin R Coombes

May 27, 2015

Contents

1	Introduction	1
2	Using SIBER	1
2.1	A Quick Example	1
2.2	Dealing With RNAseq Normalization	2
2.3	Parallelizing SIBER	3
3	Fitting Two-component Mixture Models	3
4	Session Info	5

1 Introduction

SIBER (**S**ystematic **I**dentification of **B**imodally **E**xpressed genes using **R**Naseq data) is an R package that effectively identifies bimodally expressed genes from RNAseq data based on Bimodality Index. SIBER models the RNAseq data in the finite mixture modeling framework and incorporates mechanisms for dealing with RNAseq normalization. Three types of mixture models are implemented, namely, the mixture of log normal, negative binomial or generalized poisson distribution. For completeness, we also add the normal mixture model that has been used to identify bimodal genes from microarray data.

SIBER proceeds in two steps. The first step fits a two-component mixture model. The second step calculates the Bimodality Index corresponding to the assumed mixture distribution. Four types of mixture models are implemented: log normal (LN), Negative Binomial (NB), Generalized Poisson (GP) and normal mixture (NL).

Besides identifying bimodally expressed genes, SIBER provides functionalities to fit 2-component mixture distribution from LN, NB and GP models. A degenerate case where one component becomes a point mass at zero (called 0-inflation) is also incorporated. The 0-inflated model is designed specifically to deal with the observed zero count in real RNAseq data.

2 Using SIBER

2.1 A Quick Example

Of course, we need to load the SIBER package.

```
> library(SIBER)
```

We simulate RNAseq count data from 1-component Negative Binomial distribution as below:

```
> set.seed(1000)
> N <- 100 # sample size
> G <- 200 # number of simulated genes
> # RNAseq count data simulated from NB model with mean 1000, dispersion=0.2
> Dat <- matrix(rnbinom(G*N, mu=1000, size=1/0.2), nrow=G)
```

We use the first gene for our illustration. We first fit the LN mixture model and calculate BI:

```
> SIBER(y=Dat[1, ], model='LN')

      mu1      mu2      sigma1      sigma2      pi1      delta      BI
6.2591722 6.9018463 0.3640354 0.3640354 0.1037998 1.7654165 0.5384533
```

To apply the NB model:

```
> SIBER(y=Dat[1, ], model='NB')

      mu1      mu2      sigma1      sigma2      pi1      delta
881.7878094 1292.5535708 322.5813919 472.2132984 0.7129706 0.9452263
      BI
0.4275971
```

To apply the GP model:

```
> SIBER(y=Dat[1, ], model='GP')

      mu1      mu2      sigma1      sigma2      pi1      delta
5.788304e+02 1.022990e+03 3.058969e+02 4.066634e+02 5.246266e-02 1.423620e+00
      BI
3.174076e-01
```

For the NL model, we first transform the data such that it follows normal mixture distribution.

```
> SIBER(y=log(Dat[1, ]+1), model='NL')

      mu1      mu2      sigma1      sigma2      pi1      delta      BI
6.2307179 6.8920766 0.3671304 0.3671304 0.1022049 1.8014273 0.5456841
```

Since the data is simulated from 1-component model, all of the calculated BIs are small indicating lack of bimodality.

2.2 Dealing With RNAseq Normalization

Previously, only the raw RNAseq count data is passed to `SIBER`. It is easy to incorporate RNAseq normalization in the mixture modeling. Currently, the RPKM [Mortazavi et al., 2008], TMM [Robinson et al., 2010b] and RLE [Anders and Huber, 2010] methods have been widely used to normalize RNAseq data. Once the normalization constant is estimated, i.e. using the `edgeR` package [Robinson et al., 2010a], we can easily calculate the BI after adjusting for the normalization.

In the following, we use `edgeR` package to calculate the normalization factor using TMM approach.

```
> library(edgeR)
> TMM <- calcNormFactors(Dat, method='TMM')
```

We now incorporate the TMM normalization into SIBER. We use the LN model below. The calculation with other models is similar. Note that our definition of the normalization factor differs from `edgeR` package. In our notation, $E[C_s] = d_s \mu_{c(s)}$ where C_s is the observed raw count for sample s , d_s is the normalization factor applied to sample s , $c(s) = \{1, 2\}$ denotes which of the two components sample s comes from and μ_1, μ_2 are mean parameters for the two components. Therefore, our definition of d_s maps the true expression level to the observed counts. In contrast, the normalization constant estimated by `edgeR` maps the observed counts to the estimated true expression. As a result, we need to pass the reciprocal of the normalization vector estimated by `edgeR` to SIBER.

```
> SIBER(y=Dat[1, ], d=1/TMM, model='LN')

      mu1      mu2    sigma1    sigma2      pi1    delta      BI
6.2543484 6.9027426 0.3665077 0.3665077 0.1042665 1.7691147 0.5406520
```

2.3 Parallelizing SIBER

When there are many genes to be fitted, we can easily parallel SIBER to speed up the computation. There are several ways for parallelization. Here we choose the `foreach` package for the backend. The workers are requested and registered by the `doSNOW` package.

```
library(doParallel)
cl <- makeCluster(3, type = "SOCK")
registerDoParallel(cl)
```

Note that the above command also works on Linux servers. However, it requests master nodes when run within R. For good practice, we can use `qsub` such that the computation is done in the compute nodes.

Below we illustrate how to use SIBER with parallel computation.

```
func <- function(i) {
  SIBER(y=Dat[i, ], model='LN')
}
BIinfo_LN <- foreach(i=1:nrow(Dat),
                     .combine='rbind',
                     .packages='SIBER') %dopar% {
func(i)
}
```

3 Fitting Two-component Mixture Models

SIBER package provides functions to fit three types of mixture models besides detecting bimodally expressed genes. These include: (1) 2-component mixture with equal dispersion or variance (E model); (2) 2-component mixture with unequal dispersion or variance (V model); (3) 0-inflated model. All three types of distributions are implemented.

The rule to fit 0-inflated model is that the observed percentage of count exceeds the user specified threshold. This rule overrides the model argument (E or V) when observed percentae of zero count exceeds the threshold.

First, we illustrate how to fit the E and V models. We use the simulated data from LN model. The gene we use is not 0-inflated. By default, the minimum observed percentage of zero is not achieved (`zeroPercentThr=0.2`). Hence, the 0-inflated model is disabled. In this case, the model specification will be effective.

```
> data(simDat)
> ind <- 1
> # true parameter generating the simulated data
> parList$LN[ind, ]

      mu1      mu2 sigma1 sigma2      pi1
      5.0      9.0      1.0      1.0      0.1

> # fit by E model
> fitLN(y=dataList$LN[ind, ], base=exp(1), eps=1, model='E')

      mu1      mu2      sigma1      sigma2      pi1
      4.6990876      9.1145635      0.9417622      0.9417622      0.1029831
      logLik      BIC
-2067.7614236      4156.7161167

> # fit by V model.
> fitLN(y=dataList$LN[ind, ], base=exp(1), eps=1, model='V')

      mu1      mu2      sigma1      sigma2      pi1
      4.6373527      9.1072744      0.7071008      0.9650604      0.1000980
      logLik      BIC
-2066.6375398      4159.7666664

>
```

Now we choose a gene that has zero inflation and illustrate how to fit a 0-inflated model:

```
> ind <- 5 # 0-inflated gene
> # true parameter generating the simulated data
> parList$LN[ind, ]

      mu1      mu2 sigma1 sigma2      pi1
      0.0      4.0      1.0      1.0      0.3

> # fit by E model. 0-inflated model is disabled by setting zeroPercentThr=1.
> # the result is biased.
> fitLN(y=dataList$LN[ind, ], base=exp(1), eps=1, model='E', zeroPercentThr=1)

      mu1      mu2      sigma1      sigma2      pi1
      0.05104833      4.01559275      0.78377975      0.78377975      0.30833327
      logLik      BIC
-914.36505926      1849.92338799

> # fit by 0-inflated model. 0-inflated model overrides the E model since percentage
> # of observed zero counts exceeds the threshold.
> fitLN(y=dataList$LN[ind, ], base=exp(1), eps=1, model='E', zeroPercentThr=0.2)
```

```

      mu1      mu2      sigma1      sigma2      pi1      logLik
0.0000000 3.9612703 0.0000000 0.9722536 0.3000000 -870.9626887
      BIC
1757.8203295
>

```

Here we see that when there is severe 0-inflation, fitting a E (or V) model gives biased estimate. Instead, our 0-inflated model works pretty well.

The usage of `fitNB()`, `fitGP()` is quite similar and is omitted in this manual.

4 Session Info

After all the computations, we close the connection to the workers.

```

      stopCluster(cl)
> getwd()
[1] "/tmp/Rtmp6Un1PL/Rbuild1b62606a1d7f/SIBER/vignettes"
> sessionInfo()

R version 3.2.0 Patched (2015-05-25 r68412)
Platform: x86_64-unknown-linux-gnu (64-bit)
Running under: Debian GNU/Linux 8 (jessie)

locale:
 [1] LC_CTYPE=en_US.UTF-8      LC_NUMERIC=C
 [3] LC_TIME=en_US.UTF-8      LC_COLLATE=C
 [5] LC_MONETARY=en_US.UTF-8  LC_MESSAGES=en_US.UTF-8
 [7] LC_PAPER=en_US.UTF-8     LC_NAME=C
 [9] LC_ADDRESS=C             LC_TELEPHONE=C
[11] LC_MEASUREMENT=en_US.UTF-8 LC_IDENTIFICATION=C

attached base packages:
[1] stats      graphics  grDevices  utils      datasets  methods   base

other attached packages:
[1] edgeR_3.10.2 limma_3.24.5 SIBER_1.0.0

loaded via a namespace (and not attached):
[1] mclust_5.0.1 tools_3.2.0

```

References

Simon Anders and Wolfgang Huber. Differential expression analysis for sequence count data. *Genome biol*, 11(10):R106, 2010.

Ali Mortazavi, Brian A Williams, Kenneth McCue, Lorian Schaeffer, and Barbara Wold. Mapping and quantifying mammalian transcriptomes by rna-seq. *Nature methods*, 5(7):621–628, 2008.

Mark D Robinson, Davis J McCarthy, and Gordon K Smyth. edgeR: a bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, 26(1):139–140, 2010a.

Mark D Robinson, Alicia Oshlack, et al. A scaling normalization method for differential expression analysis of rna-seq data. *Genome Biol*, 11(3):R25, 2010b.