

Differential analyses with DSS

Hao Wu

Department of Biostatistics and Bioinformatics

Emory University

Atlanta, GA 303022

`hao.wu@emory.edu`

October 15, 2013

Contents

1	Introduction	1
2	Using DSS for differential expression analysis	2
2.1	Single factor experiment	2
2.2	Multifactor experiment	4
3	Using DSS for differential methylation analysis	5
4	Session Info	6

Abstract

This vignette introduces the use of the Bioconductor package DSS (Dispersion Shrinkage for Sequencing data), which is designed for differential analysis based on high-throughput sequencing data. It performs differential expression analyses for RNA-seq, and differential methylation analyses for bisulfite sequencing (BS-seq) data. The core of DSS is a new procedure to estimate and shrink gene- or CpG site-specific dispersions, then conduct Wald tests for differential expression/methylation. Compared with existing methods, DSS provides excellent statistical and computational performance.

1 Introduction

Recent advances in high-throughput sequencing technology have revolutionized genomic research. For example, RNA-seq is a new technology for measuring the abundance of RNA products. Compared to gene expression microarrays, it provides a better dynamic range and lower signal-to-noise ratio. Bisulfite sequencing (BS-seq) is a new technology for measuring DNA methylation. Compared to capture-based methods such as MeDIP-seq, it provides single-base resolution and eliminates biases associated with CpG density.

Fundamental questions for RNA-seq or BS-seq data analyses are whether gene expression regulation or DNA methylation dynamics vary under different biological contexts. Identifying sites or regions exhibiting differential expression (DE) or differential methylation (DM) are thus key tasks in functional genomics research.

RNA- or BS-seq experiments typically have a limited number of biological replicates due to cost constraints. This can lead to unstable estimation of within group variance, and subsequently undesirable results from hypothesis testing. Variance shrinkage methods have been widely used in DE analyses based on microarray data. The methods are typically based on a Bayesian hierarchical model, with a prior imposed on the gene-specific variances to provide a basis for information sharing across all genes/CpG sites. In these models, shrinkage is achieved for variance estimation. Using shrunk variance in hypothesis tests has been shown to provide better results.

A distinct feature of RNA-seq or BS-seq data is that the measurements are in the form of counts. These data are often assumed to be from the Poisson (for RNA-seq) or Binomial (for BS-seq) distributions. Unlike continuous distributions such as the Gaussian distribution, the variances depend on means in these discrete distributions. This implies that the sample variances do not account for biological variation between replicates, and shrinkage cannot be applied on variances directly.

In contrast, we assume that our count data come from the Gamma-Poisson (RNA-seq) or Beta-Binomial (BS-seq) distribution. These distributions can be parameterized by a mean and an over dispersion parameter. The over dispersion parameters, which represent the biological variation for replicates within a treatment group, play a central role in the differential analyses.

Here we present a new DE/DM detection algorithm, where shrinkage is performed on the dispersion parameters. We first impose a log-normal prior on the dispersions, and then combine data from all genes/CpG sites to shrink dispersions through a penalized likelihood approach. Finally, we construct Wald tests to test each gene/site for differential expression/methylation. Our results show that the new method provides excellent performance compared to existing methods, especially when the overall dispersion level is high or the number of replicates is small.

Currently DSS only supports comparison of expression or methylation from two treatment groups. Methods for more advanced study designs are under development and will be implemented soon.

2 Using DSS for differential expression analysis

2.1 Single factor experiment

Required inputs for DSS are (1) gene expression values as a matrix of integers, rows are for genes and columns are for samples; and (2) a vector representing experimental designs. The length of the design vector must match the number of columns of input counts. Optionally, normalization factors or additional annotation for genes can be supplied.

The basic data container in the package is `SeqCountSet` class, which is directly inherited from `ExpressionSet` class defined in `Biobase`. An object of the class contains all necessary information for a DE analysis: gene expression values, experimental designs, and additional annotations.

A typical DE analysis contains the following simple steps.

1. Create a `SeqCountSet` object using `newSeqCountSet`.

2. Estimate normalization factor using `estNormFactors`.
3. Estimate and shrink gene-wise dispersion using `estDispersion`
4. Two-group comparison using `waldTest`.

The usage of DSS is demonstrated in the simple simulation below.

1. First load in the library, and make a `SeqCountSet` object from some counts for 2000 genes and 6 samples.

```
> library(DSS)
> counts1=matrix(rnbinom(300, mu=10, size=10), ncol=3)
> counts2=matrix(rnbinom(300, mu=50, size=10), ncol=3)
> X1=cbind(counts1, counts2) ## these are 100 DE genes
> X2=matrix(rnbinom(11400, mu=10, size=10), ncol=6)
> X=rbind(X1,X2)
> designs=c(0,0,0,1,1,1)
> seqData=newSeqCountSet(X, designs)
> seqData
```

```
SeqCountSet (storageMode: lockedEnvironment)
assayData: 2000 features, 6 samples
  element names: exprs
protocolData: none
phenoData
  sampleNames: 1 2 ... 6 (6 total)
  varLabels: designs
  varMetadata: labelDescription
featureData: none
experimentData: use 'experimentData(object)'
Annotation:
```

2. Estimate normalization factor.
3. Estimate and shrink gene-wise dispersions
4. With the normalization factors and dispersions ready, the two-group comparison can be conducted via a Wald test:

```
> result=waldTest(seqData, 0, 1)
> head(result,5)
```

	geneIndex	muA	muB	lfc	difExpr	stats	pval
38	38	4.631646	47.15068	-2.228470	-42.51903	-5.598473	2.162480e-08
91	91	8.632446	61.64730	-1.917674	-53.01486	-5.429956	5.636793e-08
53	53	9.619051	65.78331	-1.879518	-56.16425	-5.424143	5.823335e-08
5	5	7.288950	56.23838	-1.985745	-48.94943	-5.410773	6.275345e-08
35	35	6.957716	51.40795	-1.940223	-44.45023	-5.367456	7.985498e-08
		local.fdr	fdr				
38		4.241215e-06	4.241215e-06				
91		6.893638e-06	5.913825e-06				
53		7.032216e-06	5.913825e-06				
5		7.350934e-06	5.913825e-06				
35		8.660139e-06	5.913825e-06				

2.2 Multifactor experiment

DSS provides functionalities for dispersion shrinkage for multifactor experimental designs. Downstream model fitting (through genealized linear model) and hypothesis testing can be performed using other packages such as `edgeR`, with the dispersions estimated from DSS.

Below is an example, based a simple simulation, to illustrate the DE analysis of a crossed design.

1. First simulate data for a 2x2 crossed experiments. Note the counts are randomly generated.

```
> library(DSS)
> library(edgeR)
> counts=matrix(rpois(800, 10), ncol=8)
> design=data.frame(gender=c(rep("M",4), rep("F",4)), strain=rep(c("WT", "Mutant"),4))
> X=model.matrix(~gender+strain, data=design)
```

2. make `SeqCountSet`, then estimate size factors and dispersion

```
> seqData=newSeqCountSet(counts, as.data.frame(X))
> seqData=estNormFactors(seqData)
> seqData=estDispersion(seqData)
```

3. Using `edgeR`'s function to do glm model fitting, but plugging in the estimated size factors and dispersion from DSS.

```
> fit.edgeR <- glmFit(counts, X, lib.size=normalizationFactor(seqData),
+                      dispersion=dispersion(seqData))
```

4. Using `edgeR`'s function to do hypothesis testing on the second parameter of the model (gender).

```
> lrt.edgeR <- glmLRT(fit.edgeR, coef=2)
> head(lrt.edgeR$table)
```

	logFC	logCPM	LR	PValue
1	0.036816784	21.09291	8.824101e-03	0.9251594

```

2 -0.217137245 21.22977 3.127512e-01 0.5759968
3 -0.169949243 21.30170 2.119424e-01 0.6452492
4 0.002624972 21.07744 4.397775e-05 0.9947088
5 0.240133918 21.21409 4.191189e-01 0.5173770
6 0.390462143 21.07592 9.850913e-01 0.3209451

```

3 Using DSS for differential methylation analysis

For BS-seq experiments, after sequence alignment and proper processing, the BS-seq data can be summarized with following information for each C position (mostly CpG sites, with the occasional CH): chromosome number, genomic coordinate, total number of reads covering the position, and number of reads showing methylation at this position. For a sample, this information needs to be saved in a simple text file, with each row representing a CpG site.

DML detection using DSS starts from several such text files. A typical DML detection contains two simple steps. Below we will use files distributed with the package to illustrate the usage of the package.

1. Load in library. Read in text files and create **BSseq** objects for two conditions. This step requires **bsseq** Bioconductor package. **BSseq** class is defined in that package.

```

> library(DSS)
> require(bsseq)
> path <- file.path(system.file(package="DSS"), "extdata")
> dat1.1 <- read.table(file.path(path, "cond1_1.txt"), header=TRUE)
> dat1.2 <- read.table(file.path(path, "cond1_2.txt"), header=TRUE)
> dat2.1 <- read.table(file.path(path, "cond2_1.txt"), header=TRUE)
> dat2.2 <- read.table(file.path(path, "cond2_2.txt"), header=TRUE)
> BS1 <- makeBSseqData( list(dat1.1, dat1.2), paste("cond1",1:2,sep=".") )
> BS2 <- makeBSseqData( list(dat2.1, dat2.2), paste("cond2",1:2,sep=".") )
> BS1

```

```

An object of type 'BSseq' with
  28913 methylation loci
  2 samples
has not been smoothed

```

2. Detect DML by calling **callDML** function.

```

> result <- callDML(BS1, BS2)
> head(result)

```

	chr	pos	mu1	mu2	diff	diff.se	stat
1	chr18	3014740	0.8750000	0.5000000	0.37500000	NaN	NaN
2	chr18	3014769	0.0625000	0.5000000	-0.43750000	NaN	NaN
3	chr18	3014904	0.5056180	0.4705882	0.03502974	0.12361857	0.283369577

```

4 chr18 3031032 0.3421053 0.1739130 0.16819222 0.11131261 1.510989785
5 chr18 3031044 0.3466667 0.3478261 -0.00115942 0.13574931 -0.008540893
6 chr18 3031065 0.4358974 0.3750000 0.06089744 0.09967494 0.610960343

      pval      fdr
1      NaN      NaN
2      NaN      NaN
3 0.7768935 1.0000000
4 0.1307911 0.9121284
5 0.9931854 1.0000000
6 0.5412258 1.0000000

```

Currently DSS only works for detecting differentially methylated loci (DML). This means that a hypothesis test is conducted at each CpG site, and the tests are performed independently. The spatial correlations among CpG sites are ignored in the current version. For whole genome BS-seq data, it might be better to apply a smoothing technique to combine information across nearby CpG sites and improve the estimation of mean methylation levels. This is under development and will be implemented in the near future.

4 Session Info

```
> sessionInfo()
```

```
R version 3.0.2 (2013-09-25)
```

```
Platform: i386-w64-mingw32/i386 (32-bit)
```

```
locale:
```

```

[1] LC_COLLATE=C
[2] LC_CTYPE=English_United States.1252
[3] LC_MONETARY=English_United States.1252
[4] LC_NUMERIC=C
[5] LC_TIME=English_United States.1252

```

```
attached base packages:
```

```

[1] splines    parallel  stats      graphics  grDevices  utils      datasets
[8] methods    base

```

```
other attached packages:
```

```

[1] bsseq_0.10.0      matrixStats_0.8.12  GenomicRanges_1.14.0
[4] XVector_0.2.0     IRanges_1.20.0      edgeR_3.4.0
[7] limma_3.18.0      DSS_1.8.0           locfdr_1.1-7
[10] Biobase_2.22.0    BiocGenerics_0.8.0

```

```
loaded via a namespace (and not attached):
```

```
[1] R.methodsS3_1.5.2 RColorBrewer_1.0-5 colorspace_1.2-4  dichromat_2.0-0
```

[5]	grid_3.0.2	labeling_0.2	lattice_0.20-24	locfit_1.5-9.1
[9]	munSELL_0.4.2	plyr_1.8	scales_0.2.3	stats4_3.0.2
[13]	stringr_0.6.2	tools_3.0.2		