

Gene set enrichment analysis of RNA-Seq data with the *SeqGSEA* package

Xi Wang^{1,2} and Murray Cairns^{1,2,3}

October 29, 2024

¹School of Biomedical Sciences and Pharmacy, The University of Newcastle, Callaghan, New South Wales, Australia

²Hunter Medical Research Institute, New Lambton, New South Wales, Australia

³Schizophrenia Research Institute, Sydney, New South Wales, Australia

`xi.wang@mdc-berlin.de`

1	Introduction	2
1.1	Background	2
1.2	Getting started	2
1.3	Package citation	2
2	Differential splicing analysis and DS scores	2
2.1	The <i>ReadCountSet</i> class	2
2.2	DS analysis and DS scores	3
2.3	DS permutation p-values	5
3	Differential expression analysis and DE scores	6
3.1	Gene read count data from <i>ReadCountSet</i> class	6
3.2	DE analysis and DE scores	6
3.3	DE permutation p-values	7
4	Integrative GSEA runs	8
4.1	DE/DS score integration	8
4.2	Initialization of <i>SeqGeneSet</i> objects	10
4.3	running GSEA with integrated gene scores	11
4.4	<i>SeqGSEA</i> result displays	11
5	Running <i>SeqGSEA</i> with multiple cores	13
5.1	R-parallel packages	13
5.2	Parallelizing analysis on permutation data sets	14
6	Analysis examples	14
6.1	Starting from your own RNA-Seq data	14
6.2	Exemplified pipeline for integrating DE and DS	15
6.3	Exemplified pipeline for DE-only analysis	17
6.4	One-step <i>SeqGSEA</i> analysis	18
7	Session information	19

1 Introduction

1.1 Background

Transcriptome sequencing (RNA-Seq) has become a key technology in transcriptome studies because it can quantify overall expression levels and the degree of alternative splicing for each gene simultaneously. Many methods and tools, including quite a few R/Bioconductor packages, have been developed to deal with RNA-Seq data for differential expression analysis and thereafter functional analysis aiming at novel biological and biomedical discoveries. However, those tools mainly focus on each gene's overall expression and may miss the opportunities for discoveries regarding alternative splicing or the combination of the two.

SeqGSEA is novel R/Bioconductor package to derive biological insight by integrating differential expression (DE) and differential splicing (DS) from RNA-Seq data with functional gene set analysis. Due to the digital feature of RNA-Seq count data, the package utilizes negative binomial distributions for statistical modeling to first score differential expression and splicing in each gene, respectively. Then, integration strategies are applied to combine the two scores for integrated gene set enrichment analysis. See the publications Wang and Cairns (2013) and Wang and Cairns (2014) for more details. The *SeqGSEA* package can also give detection results of differentially expressed genes and differentially spliced genes based on sample label permutation.

1.2 Getting started

The *SeqGSEA* depends on *Biobase* for definitions of class *ReadCountSet* and class *SeqGeneSet*, *DESeq2* for differential expression analysis, *biomaRt* for gene IDs/names conversion, and *doParallel* for parallelizing jobs to reduce running time. Make sure you have these dependent packages installed before you install *SeqGSEA*.

To load the *SeqGSEA* package, type `library(SeqGSEA)`. To get an overview of this package, type `?SeqGSEA`.

```
> library(SeqGSEA)
```

```
> ? SeqGSEA
```

In this Users' Guide of the *SeqGSEA* package, an analysis example is given in Section 6, and detailed guides for DE, DS, and integrative GSEA analysis are given in Sections 3, 2, and 4, respectively. A guide to parallelize those analyses is given in Section 5.

1.3 Package citation

To cite this package, please cite the article below:

Wang X and Cairns MJ (2014). **SeqGSEA: a Bioconductor package for gene set enrichment analysis of RNA-Seq data integrating differential expression and splicing.** *Bioinformatics*, **30**(12):1777-9.

To cite/discuss the method used in this package, please cite the article below:

Wang X and Cairns MJ (2013). **Gene set enrichment analysis of RNA-Seq data: integrating differential expression and splicing.** *BMC Bioinformatics*, **14**(Suppl 5):S16.

2 Differential splicing analysis and DS scores

2.1 The *ReadCountSet* class

To facilitate differential splicing (DS) analysis, *SeqGSEA* saves exon read count data using *ReadCountSet* class, which is derived from *eSet*. While below is an example showing the steps to create a new *ReadCountSet* object, creating a *ReadCountSet* object from your own data should refer to Section 6.

```

> rcounts <- cbind(t(sapply(1:10, function(x) {rnbinom(5, size=10, prob=runif(1))})),
+                 t(sapply(1:10, function(x) {rnbinom(5, size=10, prob=runif(1))})))
> colnames(rcounts) <- c(paste("S", 1:5, sep=""), paste("C", 1:5, sep=""))
> geneIDs <- c(rep("G1", 4), rep("G2", 6))
> exonIDs <- c(paste("E", 1:4, sep=""), paste("E", 1:6, sep=""))
> RCS <- newReadCountSet(rcounts, exonIDs, geneIDs)
> RCS

```

```

ReadCountSet (storageMode: environment)
assayData: 10 features, 10 samples
  element names: counts
protocolData: none
phenoData
  sampleNames: S1 S2 ... C5 (10 total)
  varLabels: label
  varMetadata: labelDescription
featureData
  featureNames: 1 2 ... 10 (10 total)
  fvarLabels: exonIDs geneIDs ... padjust (10 total)
  fvarMetadata: labelDescription
experimentData: use 'experimentData(object)'
Annotation:

```

2.2 DS analysis and DS scores

To better illustrate DS analysis functions, we load an example `ReadCountSet` object from a real RNA-Seq data set as follows.

```

> data(RCS_example, package="SeqGSEA")
> RCS_example

ReadCountSet (storageMode: environment)
assayData: 5000 features, 20 samples
  element names: counts
protocolData: none
phenoData
  sampleNames: S1 S2 ... C10 (20 total)
  varLabels: label
  varMetadata: labelDescription
featureData
  featureNames: ENSG00000000003:001 ENSG00000000003:002 ...
  ENSG00000007402:038 (5000 total)
  fvarLabels: exonIDs geneIDs ... padjust (10 total)
  fvarMetadata: labelDescription
experimentData: use 'experimentData(object)'
Annotation:

```

This example `ReadCountSet` object is comprised of 20 samples and 5,000 exons, part of the prostate cancer RNA-Seq data set (Kannan et al., 2011). With the function `geneID` and the script below, we can easily check the number of genes involved in this data set.

```

> length(unique(geneID(RCS_example)))

[1] 182

```

Noticed that some exons are too short or not expressed, we should first filter out these exons from following analysis to secure the robustness of our analysis. By default, function `exonTestability` marks exons with the sum of read counts across all samples less than `cutoff` (default: 5) to be excluded in downstream analysis. Users can also exclude genes with no or low expression from downstream analysis by checking `geneTestability`.

```
> RCS_example <- exonTestability(RCS_example, cutoff = 5)
```

Then, the main DS analysis is executed using function `estiExonNBstat` for exon DS NB-statistics and function `estiGeneNBstat` for gene DS NB-statistics by averaging exon NB-statistics. Please refer to Wang et al. (2013) for detailed statistic analysis regarding differential splicing from exon count data.

```
> RCS_example <- estiExonNBstat(RCS_example)
> RCS_example <- estiGeneNBstat(RCS_example)
> head(fData(RCS_example)[, c("exonIDs", "geneIDs", "testable", "NBstat")])
```

	exonIDs	geneIDs	testable	NBstat
ENSG00000000003:001	E001	ENSG00000000003	TRUE	2.0219857
ENSG00000000003:002	E002	ENSG00000000003	TRUE	0.2486443
ENSG00000000003:003	E003	ENSG00000000003	TRUE	0.1238136
ENSG00000000003:004	E004	ENSG00000000003	TRUE	1.2058520
ENSG00000000003:005	E005	ENSG00000000003	TRUE	2.0668287
ENSG00000000003:006	E006	ENSG00000000003	TRUE	0.2678247

We run DS analysis on the permutation data sets as well. Here we set to run permutation 20 times for demonstration; however, in practice at least 1,000 permutations are recommended. To do so, we first generate a permutation matrix, each column corresponding to each permutation; then run DS analysis on the permutation data sets, and updated `permute_NBstat_gene` slot for results.

```
> permuteMat <- genpermuteMat(RCS_example, times=20)
> RCS_example <- DSpermute4GSEA(RCS_example, permuteMat)
> head(RCS_example@permute_NBstat_gene)
```

	result.1	result.2	result.3	result.4	result.5	result.6
ENSG00000000003	0.6271838	0.9904407	1.5281199	1.0415207	0.6534197	3.1995316
ENSG00000000005	1.1865299	0.5560025	0.7398135	0.1935409	0.3875388	0.2769597
ENSG00000000419	1.3266792	0.6293865	0.5732224	0.5425767	0.3729265	1.9342743
ENSG00000000457	1.3277575	1.7571730	0.6756801	0.6529276	0.7872754	1.3379359
ENSG00000000460	0.8562294	0.8512607	0.5650345	0.7950802	0.6987681	0.8215001
ENSG00000000938	1.2032345	1.2153165	1.1187942	1.2906723	1.1799358	1.2335696
	result.7	result.8	result.9	result.10	result.11	result.12
ENSG00000000003	1.0278354	1.7736611	3.245086	1.1158553	1.752466	0.7929791
ENSG00000000005	0.5835158	0.2146176	1.045034	0.5081923	0.359773	0.4993746
ENSG00000000419	0.9918119	0.4725558	4.681872	0.6785681	2.160443	1.1609020
ENSG00000000457	0.6972578	1.7138230	2.102344	0.6917472	1.878783	1.5445549
ENSG00000000460	0.5451070	0.9707300	1.063114	0.6344656	1.555101	0.7846695
ENSG00000000938	1.2436998	1.6499414	1.766332	1.6420525	1.331123	1.9144804
	result.13	result.14	result.15	result.16	result.17	result.18
ENSG00000000003	1.3519782	1.1801358	0.6980744	0.8766650	0.2888305	1.3726854
ENSG00000000005	0.3425306	0.4959407	0.5324158	0.2827720	0.1328424	0.5443518
ENSG00000000419	0.8976075	0.8413172	0.6571168	0.2831129	0.4512044	1.0058709
ENSG00000000457	0.8137042	0.4890323	0.5460917	0.5278108	1.1897145	1.8898783
ENSG00000000460	0.6088692	0.7027171	0.4498368	0.5250597	1.1124061	0.7736096
ENSG00000000938	1.0514047	1.0549320	1.2951907	1.0257995	0.8942183	1.5877153

```

                result.19 result.20
ENSG000000000003 0.4057037 0.4652395
ENSG000000000005 0.5841833 0.5581997
ENSG000000000419 0.5579777 0.3188484
ENSG000000000457 0.9465964 1.0369771
ENSG000000000460 0.4979725 0.8273467
ENSG000000000938 1.5787000 1.3562525

```

The DS NB-statistics from the permutation data sets offer an empirical background of NB-statistics on the real data set. By normalizing NB-statistics against this background, we get the DS scores, which will be used in integrated GSEA runs (Section 4).

```

> DSscore.normFac <- normFactor(RCS_example@permute_NBstat_gene)
> DSscore <- scoreNormalization(RCS_example@featureData_gene$NBstat,
+                               DSscore.normFac)
> DSscore.perm <- scoreNormalization(RCS_example@permute_NBstat_gene,
+                                    DSscore.normFac)
> DSscore[1:5]

ENSG000000000003 ENSG000000000005 ENSG000000000419 ENSG000000000457 ENSG000000000460
      1.1822528      0.9107181      1.6398058      2.5738879      0.9605115

> DSscore.perm[1:5,1:10]

                result.1 result.2 result.3 result.4 result.5 result.6
ENSG000000000003 0.5143505 0.8122557 1.2532038 0.8541461 0.5358664 2.6239206
ENSG000000000005 2.3673475 1.1093282 1.4760654 0.3861500 0.7732120 0.5525861
ENSG000000000419 1.2919092 0.6128913 0.5581992 0.5283567 0.3631527 1.8835802
ENSG000000000457 1.1746395 1.5545344 0.5977601 0.5776314 0.6964861 1.1836440
ENSG000000000460 1.0950010 1.0886468 0.7226023 1.0167995 0.8936295 1.0505870
                result.7 result.8 result.9 result.10
ENSG000000000003 0.8429229 1.4545710 2.661280 0.9151076
ENSG000000000005 1.1642224 0.4282020 2.085038 1.0139380
ENSG000000000419 0.9658182 0.4601709 4.559168 0.6607840
ENSG000000000457 0.6168495 1.5161835 1.859900 0.6119743
ENSG000000000460 0.6971177 1.2414318 1.359578 0.8113953

```

2.3 DS permutation p-values

Besides calculating DS scores, based on the NB statistics on the real data set and the permutation data sets, we can also calculate a permutation p-value for each gene's DS significance in the studied data set.

```

> RCS_example <- DSpermutePval(RCS_example, permuteMat)
> head(DSresultGeneTable(RCS_example))

      geneID  NBstat pvalue  padjust
1 ENSG000000000003 1.4416043  0.25 0.3697479
2 ENSG000000000005 0.4564578  0.60 0.6478528
3 ENSG000000000419 1.6839390  0.15 0.2838710
4 ENSG000000000457 2.9094026  0.00 0.0000000
5 ENSG000000000460 0.7510661  0.55 0.6165605
6 ENSG000000000938 1.7949049  0.05 0.1466667

```

The adjusted p-values accounting for multiple testings were given by the BH method (Benjamini and Hochberg, 1995). Users can also apply function `topDSGenes` and function `topDSExons` to quickly get the most significant DS genes and exons, respectively.

3 Differential expression analysis and DE scores

3.1 Gene read count data from *ReadCountSet* class

For gene DE analysis, read counts on each gene should be first calculated. With *SeqGSEA*, users usually analyze DE and DS simultaneously, so the package includes the function `getGeneCount` to facilitate gene read count calculation from a *ReadCountSet* object.

```
> geneCounts <- getGeneCount(RCS_example)
> dim(geneCounts) # 182 20

[1] 182 20

> head(geneCounts)

           S1  S2  S3  S4  S5  S6  S7  S8  S9 S10 C1  C2  C3  C4  C5
ENSG000000000003 495 235 386 272 255 815 1065 803 839 885 278 270 238 175 292
ENSG000000000005  19   1   0   2   2  12   7   3   1   4   4   4   2   0   1
ENSG000000000419 196 134 165 184 132 344  343 307 342 280 179 156 100 120 126
ENSG000000000457  97  78 141  72 102 219  344 277 337 249  62  48  40  43  52
ENSG000000000460  52  35  48  25  47 105  124  80 156 145  48  36  21  34  19
ENSG000000000938  27  44  57  43  14  71   74 146 148 165  48  59  32  79  20
           C6  C7  C8  C9 C10
ENSG000000000003 432 519 621 475 560
ENSG000000000005  9   3   1  14  46
ENSG000000000419 169 255 171 164 201
ENSG000000000457 170 165 131 183 185
ENSG000000000460  68  90  48  72  38
ENSG000000000938 103 1285 137 156  90
```

This function results in a matrix of 182 rows and 20 columns, corresponding to 182 genes and 20 samples.

3.2 DE analysis and DE scores

DE analysis has been implemented in several R/Bioconductor packages, of which *DESeq2* (Anders and Huber, 2010) is mainly utilized in *SeqGSEA* for DE analysis. With *DESeq2*, we can model count data with negative binomial distributions for accounting biological variations and various biases introduced in RNA-Seq. Given the read count data on individual genes and sample grouping information, basic DE analysis based on *DESeq2* including size factor estimation and dispersion estimation, is encapsulated in the function `runDESeq`.

```
> label <- label(RCS_example)
> dds <- runDESeq(geneCounts, label)
```

The function `runDESeq` returns a *DESeqDataSet* object, which is defined in the *DESeq2* package. The DE analysis in the *DESeq2* package continues with the output *DESeqDataSet* object and conducts negative-binomial-based statistical tests for DE genes (using `nbinomTest` or `nbinomGLMTest`). However, in this *SeqGSEA* package, we define NB statistics to quantify each gene's expression difference between sample groups.

The NB statistics for DE can be achieved by the following scripts.

```
> DEGres <- DENBStat4GSEA(dds)
> DEGres[1:5, "NBstat"]

[1] 0.51449518 0.32697507 0.02205611 13.22540371 1.24725584
```

Similarly, we run DE analysis on the permutation data sets as well. The `permuteMat` should be the same as used in DS analysis on the permutation data sets.

```
> DEpermNBstat <- DENBStatPermut4GSEA(dds, permuteMat)
> DEpermNBstat[1:5, 1:10]

      result.1 result.2 result.3 result.4 result.5 result.6
[1,] 0.777051082 0.94538875 2.11173190 0.01843755 1.042230463 2.349866232
[2,] 0.161417438 0.06037666 0.07973917 3.14746533 3.386609705 0.404190515
[3,] 1.401502791 1.19754393 1.66313269 0.37570495 0.002763602 3.856454623
[4,] 0.056857708 0.42152012 0.06732825 0.11388675 0.846310005 0.277667363
[5,] 0.005528121 0.10208625 0.25868381 0.74997883 0.880428877 0.002557039
      result.7 result.8 result.9 result.10
[1,] 0.72917183 0.7201341 4.90724969 3.3179554
[2,] 0.03528390 1.0947574 0.05552578 1.5946220
[3,] 2.00264563 2.1661580 4.34540909 0.5750868
[4,] 0.00287456 0.9230497 1.03226627 1.2774004
[5,] 0.09366664 0.7292835 0.22618874 1.5568079
```

Once again, the DE NB-statistics from the permutation data sets offer an empirical background, so we can normalize NB-statistics against this background. By doing so, we get the DE scores, which will also be used in integrated GSEA runs (Section 4).

```
> DEScore.normFac <- normFactor(DEpermNBstat)
> DEScore <- scoreNormalization(DEGres$NBstat, DEScore.normFac)
> DEScore.perm <- scoreNormalization(DEpermNBstat, DEScore.normFac)
> DEScore[1:5]
```

```
[1] 0.29261115 0.31369757 0.01504121 12.34586005 1.90097052
```

```
> DEScore.perm[1:5, 1:10]
```

```
      result.1 result.2 result.3 result.4 result.5 result.6
[1,] 0.441935746 0.53767518 1.20101475 0.01048607 0.592752404 1.336449958
[2,] 0.154862745 0.05792494 0.07650119 3.01965591 3.249089328 0.387777513
[3,] 0.955757654 0.81666750 1.13417669 0.25621275 0.001884644 2.629917005
[4,] 0.053076437 0.39348730 0.06285064 0.10631281 0.790026914 0.259201343
[5,] 0.008425532 0.15559194 0.39426578 1.14305951 1.341881345 0.003897241
      result.7 result.8 result.9 result.10
[1,] 0.41470517 0.4095651 2.79092212 1.8870357
[2,] 0.03385113 1.0503025 0.05327104 1.5298691
[3,] 1.36570823 1.4772158 2.96336049 0.3921816
[4,] 0.00268339 0.8616631 0.96361633 1.1924480
[5,] 0.14275942 1.1115172 0.34473932 2.3727657
```

3.3 DE permutation p-values

Similar to DS analysis, comparing NB-statistics on the real data set and those on the permutation data sets, we can get permutation p-values for each gene's DE significance.

```
> DEGres <- DEpermutePval(DEGres, DEpermNBstat)
> DEGres[1:6, c("NBstat", "perm.pval", "perm.padj")]
```

```
      NBstat perm.pval perm.padj
ENSG00000000003 0.51449518 0.75 1
```

```

ENSG000000000005  0.32697507      0.60      1
ENSG000000000419  0.02205611      0.85      1
ENSG000000000457 13.22540371      0.00      0
ENSG000000000460  1.24725584      0.15      1
ENSG000000000938  7.73813817      0.00      0

```

For a comparison to the nominal p-values from exact testing and forming comprehensive results, users can run `DENBTest` first and then `DEpermutePval`, which generates results as follows.

```

> DEGres <- DENBTest(dds)
> DEGres <- DEpermutePval(DEGres, DEpermNBstat)
> DEGres[1:6, c("NBstat", "pval", "padj", "perm.pval", "perm.padj")]

```

	NBstat	pval	padj	perm.pval	perm.padj
ENSG000000000003	0.51449518	0.1112841241	0.24121931	0.75	1
ENSG000000000005	0.32697507	0.2310245778	0.38574746	0.60	1
ENSG000000000419	0.02205611	0.3643336232	0.52272077	0.85	1
ENSG000000000457	13.22540371	0.0008268416	0.01588092	0.00	0
ENSG000000000460	1.24725584	0.0816750168	0.21235504	0.15	1
ENSG000000000938	7.73813817	0.0006694334	0.01588092	0.00	0

4 Integrative GSEA runs

4.1 DE/DS score integration

We have proposed two strategies for integrating normalized DE and DS scores (Wang and Cairns, 2013), one of which is the weighted summation of the two scores and the other is a rank-based strategy. The functions `geneScore` and `genePermuteScore` implement two methods for the weighted summation strategy: weighted linear combination and weighted quadratic combination. Scripts below show a linear combination of DE and DS scores with weight for DE equal to 0.3. Users should keep the weight for DE in `geneScore` and `genePermuteScore` the same, and the weight ranges from 0 (i.e., DS only) to 1 (i.e., DE only). Visualization of gene scores can be made by applying the `plotGeneScore` function.

```

> gene.score <- geneScore(DEscore, DSscore, method="linear", DEweight = 0.3)
> gene.score.perm <- genePermuteScore(DEscore.perm, DSscore.perm,
+                                   method="linear", DEweight=0.3)
> plotGeneScore(gene.score, gene.score.perm)

```

The plot generated by the `plotGeneScore` function (Fig. 1) can also be saved as a PDF file easily with the `pdf` argument of `plotGeneScore`.

The functions `geneScore` and `genePermuteScore` also implement one method for the rank-based integration strategy: using data-set-specific ranks. The plot for integrated gene scores is shown in Fig. 2.

```

> gene.score <- geneScore(DEscore, DSscore, method="rank", DEweight = 0.3)
> gene.score.perm <- genePermuteScore(DEscore.perm, DSscore.perm,
+                                   method="rank", DEweight=0.3)
> plotGeneScore(gene.score, gene.score.perm)

```

Rather than the above method to integrate scores with data-set-specific ranks, an alternative method is implemented with the `rankCombine` function, which takes only the ranks from the real data set for integrating DE and DS scores on both real and permutation data sets. This provides a method in a global manner. The plot of gene scores is shown in Fig. 3.




Figure 1: Gene scores resulted from linear combination. Scores are sorted from the largest to the smallest. Red, green, orange, blue dotted horizontal lines represent the maximum score, average score on the real data set, and the maximum score, average score on the permutation data sets.




Figure 2: Gene scores resulted from rank-based combination with data-set-specific ranks. Scores are sorted from the largest to the smallest. Red, green, orange, blue dotted horizontal lines represent the maximum score, average score on the real data set, and the maximum score, average score on the permutation data sets.




Figure 3: Gene scores resulted from rank-based combination with the same rank got from the real data set. Scores are sorted from the largest to the smallest. Red, green, orange, blue dotted horizontal lines represent the maximum score, average score on the real data set, and the maximum score, average score on the permutation data sets.

```
> combine <- rankCombine(DEscore, DSscore, DEscore.perm, DSscore.perm, DEweight=0.3)
> gene.score <- combine$geneScore
> gene.score.perm <- combine$genePermuteScore
> plotGeneScore(gene.score, gene.score.perm)
```

Basically the integrated gene scores are distributed similarly with the three integration methods at DE weight 0.3 (Figs. 1, 2, and 3); however, according to the analysis in Wang and Cairns (2013), SeqGSEA can detect slightly more significant gene sets with rank-based integration strategy than with linear combination.

4.2 Initialization of *SeqGeneSet* objects

To facilitate running gene set enrichment analysis, *SeqGSEA* implements a *SeqGeneSet* class. The *SeqGeneSet* class has several slots for accommodating a category of gene sets derived from any biological knowledge-based databases such as Kyoto Encyclopedia of Genes and Genomes (KEGG). However, we recommend to start with the formatted gene-set files from the well-maintained resource Molecular Signatures Database (MSigDB, <http://www.broadinstitute.org/gsea/msigdb/index.jsp>) (Subramanian et al., 2005). After downloading a gmt file from the above URL, users can use `loadGenesets` to initialize a *SeqGeneSet* object easily. Please note that with the current version of *SeqGSEA*, only gene sets with gene symbols are supported, though read count data's gene IDs can be either gene symbols or Ensembl Gene IDs.

Below is shown an example of the *SeqGeneSet* object, which contains information such as how many gene sets in this object and the names/sizes/descriptions of each gene set.

```
> data(GS_example, package="SeqGSEA")
> GS_example
```

```

SeqGeneSet object: gs_symb.txt
GeneSetSourceFile: /Library/Frameworks/R.framework/Versions/2.15/Resources/library/SeqGSEA/extdat
GeneSets: ERB2_UP.V1_DN
          AKT_UP_MTOR_DN.V1_UP
          ...
          KRAS.600.LUNG.BREAST_UP.V1_DN
with the number of genes in respective sets: 6, 6, ..., 5
brief descriptions:
  http://www.broadinstitute.org/gsea/msigdb/cards/ERB2_UP.V1_DN
  http://www.broadinstitute.org/gsea/msigdb/cards/AKT_UP_MTOR_DN.V1_UP
  ...
  http://www.broadinstitute.org/gsea/msigdb/cards/KRAS.600.LUNG.BREAST_UP.V1_DN
# gene sets passed filter: 11 (#genes >= 5 AND <= 1000)
# gene sets excluded: 178 (#genes < 5 OR > 1000)
ES scores: not computed
ES postions: not computed
Permutated ES scores: not performed
ES scores normalized: No
ES p-value: not computed
ES FWER: not computed
ES FDR: not computed

```

4.3 running GSEA with integrated gene scores

With the initialized `SeqGeneSet` object and integrated gene scores as well as gene scores on the permutation data sets, the main `GSEnrichAnalyze` can be executed; and the `topGeneSets` allows users promptly access to the top significant gene sets.

```

> GS_example <- GSEnrichAnalyze(GS_example, gene.score, gene.score.perm)
> topGeneSets(GS_example, 5)

```

	GSName	GSSize	ES	ES.pos	pval	FDR	FWER
8	HOXA9_DN.V1_UP	5	2.0777382	25	0.05	0.00000	0.05
9	TBK1.DF_UP	5	2.0507110	17	0.00	0.00000	0.05
7	BMI1_DN.V1_DN	6	1.4538528	12	0.20	0.25000	0.80
11	KRAS.600.LUNG.BREAST_UP.V1_DN	5	1.4888687	52	0.10	0.33333	0.75
6	PTEN_DN.V2_DN	5	0.8628552	106	0.65	0.87500	1.00

The main GSEA includes several steps detailed in [Wang and Cairns \(2013\)](#) and its original paper [Subramanian et al. \(2005\)](#). In *SeqGSEA*, functions `caLES`, `caLES.perm`, `normES` and `signifES` are implemented to complete the analysis. Advanced users may set up customized pipelines with the functions above themselves.

4.4 SeqGSEA result displays

Several functions in *SeqGSEA* can be employed for visualization of gene set enrichment analysis running results. The `plotES` function is to plot the distribution of normalized enrichment scores (NES) of all gene sets in a `SeqGeneSet` object on the observed data set versus its empirical background provided by the NES on the permutation data sets (Fig. 4).

```

> plotES(GS_example)

```

The `plotSig` function plots the distributions of permutation p -value, false discovery rate (FDR) and family-wise error rate (FWER) versus NES. The example plot is not shown in this vignette as the distributions can be far from the real ones due to the limited permutation times.

Global Observed and Null Densities (Area Normalized)




Figure 4: Distribution of normalized enrichment scores (NES) on the observed and permutation (null) data sets. Blue: observed NES density; Orange: each for NES density on one permutation data set; Red: the average density on all permutation data sets; Black: observed NES values.

```
> plotSig(GS_example)
```

The `plotSigGS` function is to plot detailed results of a particular gene set that has been analyzed. Information in the plot includes running enrichment scores, null NES on the permutation data sets. See Fig. 5 for an example.

```
> plotSigGeneSet(GS_example, 9, gene.score) # 9th gene set is the most significant one.
```

Besides the functions to generate plots, the `writeSigGeneSet` function can write the detailed information of any analyzed gene sets, including NES, p-values, FDR, and the leading set (see the definition in Wang and Cairns (2013)). An example is shown below.

```
> writeSigGeneSet(GS_example, 9, gene.score) # 9th gene set is the most significant one.
```

GSEA result for gene set No. 9:

```
genesetName      gs_symb.txt:TBK1.DF_UP
genesetSize      5
genesetDesc      http://www.broadinstitute.org/gsea/msigdb/cards/TBK1.DF_UP
NES              2.0507109942745
Pos              17
pvalue           0
FDR              0
FWER             0.05
```

Leading set:

```
ENSG00000005194      6.18977240025557
ENSG00000001167      4.46759112115815
ENSG00000002919      3.99165269533001
```

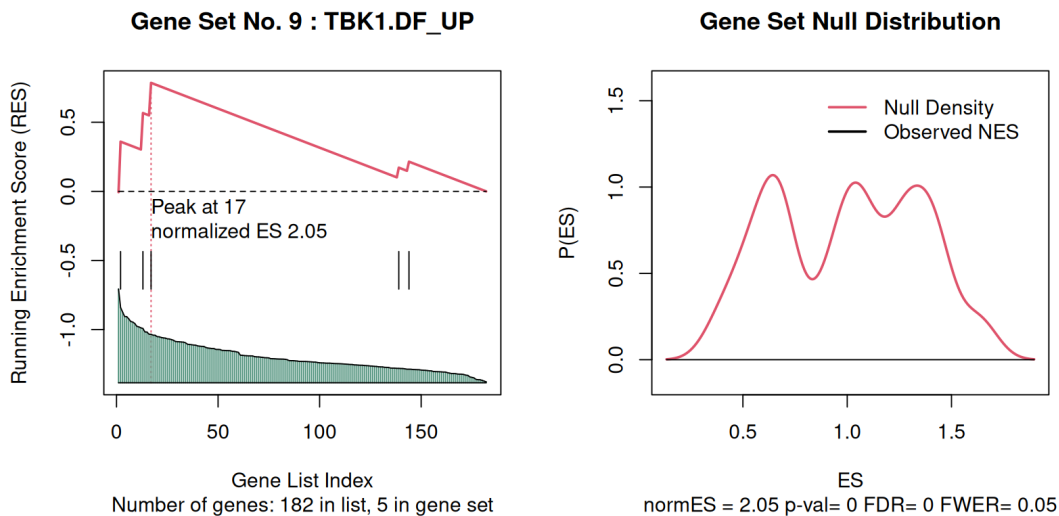


Figure 5: Left: gene locations of a particular gene set according to the gene score rank and running enrichment scores; Right: null NES distribution and the relative position of the observed NES.

Whole gene set:

ENSG00000005194	6.18977240025557
ENSG00000001167	4.46759112115815
ENSG00000002919	3.99165269533001
ENSG00000006576	1.18711620122735
ENSG00000005059	1.11723300635901

The `GSEAResultTable` generates a summary table of the GSEA analysis, which can also be output with customized scripts. An example can be found in Section 6.

5 Running *SeqGSEA* with multiple cores

5.1 R-parallel packages

There are many R packages for facilitating users in running R scripts in parallel, including *parallel*, *snowfall*, *multicore*, and many others. While experienced users may parallelize *SeqGSEA* runnings with the above packages themselves to reduce the running time, we provide with in the *SeqGSEA* package vignette an general way for users to parallelize their runnings utilizing the *doParallel* package (which depends on *parallel*).

First, we show a toy example for a basic idea how *doParallel* works. Basically, *doParallel* is a *parallel backend* for the *foreach* package using *parallel*, which provides a mechanism to execute *foreach* loops in parallel. With the `foreach` function in the *foreach* package, we can specify which `foreach` loops need to be parallelized using the `%dopar%` operator. However, without a registered parallel backend, the `foreach` loops will be executed sequentially even if the `%dopar%` operator is used. In those cases, the *foreach* package will issue a warning that it is running sequentially. Below are two running examples showing how the task is running sequentially and in parallel, respectively.

Run sequentially without parallel backend registered

```
> library(doParallel)
> a <- matrix(1:16, 4, 4)
```

```

> b <- t(a)
> foreach(b=iter(b, by='col'), .combine=cbind) %dopar%
+   (a %**% b)

      [,1] [,2] [,3] [,4]
[1,] 276 304 332 360
[2,] 304 336 368 400
[3,] 332 368 404 440
[4,] 360 400 440 480

```

Although the warning message didn't appear here, you would definitely see a warning message when you run the scripts above, like:

Warning message:

executing %dopar% sequentially: no parallel backend registered

Run in parallel with two cores

```

> library(doParallel)
> cl <- makeCluster(2) # specify 2 cores to be used in this computing
> registerDoParallel(cl)
> getDoParWorkers() # 2

[1] 2

> a <- matrix(1:16, 4, 4)
> b <- t(a)
> foreach(b=iter(b, by='col'), .combine=cbind) %dopar%
+   (a %**% b)

      [,1] [,2] [,3] [,4]
[1,] 276 304 332 360
[2,] 304 336 368 400
[3,] 332 368 404 440
[4,] 360 400 440 480

```

The parallel backend registration was done with `registerdoParallel`. For more details please refer to `doParallel`'s vignette (<http://cran.r-project.org/web/packages/doParallel/index.html>).

5.2 Parallelizing analysis on permutation data sets

In *SeqGSEA*, the loops for analyzing permutation data sets are implemented by `foreach` with `%dopar%` operator used. Those loops include DS, DE, and GSEA analyses, which are the most time consuming parts. Although there are three parts can take the advantage of parallel running, users only need to register parallel backend once at the beginning of all analyses. See an analysis example in the next section (Section 6).

6 Analysis examples

6.1 Starting from your own RNA-Seq data

With this *SeqGSEA* package, we provide complementary Python scripts for counting reads on exons of each genes from SAM/BAM files: two scripts `prepare_exon_annotation_refseq.py` and `prepare_exon_annotation_ensembl.py` for preparing (sub-)exon annotation, and `count_in_exons.py` for counting reads. The scripts are based on the HTSeq Python package (<http://www-huber.embl.de/users/anders/HTSeq/>). Please install it before using the Python scripts provided. The scripts can be found in the directory given by the following command.

```
> system.file("extscripts", package="SeqGSEA", mustWork=TRUE)
```

```
[1] "/tmp/RtmpuycajV/Rinst1b32ae41074064/SeqGSEA/extscripts"
```

Simply by typing “python” + the file name of echo script in your shell console, the help documentation will be on your screen.

Other than the Python scripts provided, users who prefer playing with R/Bioconductor packages can also use `easyRNASeq` in *easyRNASeq*, `summarizeOverlaps` in *GenomicRanges*, and `featureCounts` in *Rsubread* to count reads that mapped to each exon. Please refer to respective packages for detailed usage.

For users who are not familiar with RNA-Seq data processing, the upstream steps of counting reads are (1) data preprocessing, including adapter removal, low-quality read filtering, data quality-control analysis, and (2) read mapping. R/Bioconductor users can apply *Rsubread* to map reads based on a seed-and-vote approach, as well as a few QC analysis. Users familiar with command-line can choose from a wide range of tools, such as already widely used ones including TopHat (<http://tophat.cbcb.umd.edu>), STAR (<http://code.google.com/p/rna-star>), and etc..

6.2 Exemplified pipeline for integrating DE and DS

Below is shown a typical SeqGSEA running example with the data enclosed with the *SeqGSEA* package, which are a part of the prostate cancer data set (Kamman et al., 2011). We divide the process into five steps for a complete SeqGSEA run.

Step 0: Initialization. (Users should change values in this part accordingly.)

```
> rm(list=ls())
> # input count data files
> data.dir <- system.file("extdata", package="SeqGSEA", mustWork=TRUE)
> case.pattern <- "~SC" # file name starting with "SC"
> ctrl.pattern <- "~SN" # file name starting with "SN"
> case.files <- dir(data.dir, pattern=case.pattern, full.names = TRUE)
> control.files <- dir(data.dir, pattern=ctrl.pattern, full.names = TRUE)
> # gene set file
> geneset.file <- system.file("extdata", "gs_symb.txt",
+                             package="SeqGSEA", mustWork=TRUE)
> # output file prefix
> output.prefix <- "SeqGSEA.test"
> # setup parallel backend
> library(doParallel)
> cl <- makeCluster(2) # specify 2 cores to be used in computing
> registerDoParallel(cl) # parallel backend registration
> # setup permutation times
> perm.times <- 20 # change the number to >= 1000 in your analysis
```

Step 1: DS analysis

```
> # load exon read count data
> RCS <- loadExonCountData(case.files, control.files)
> # remove genes with low expression
> RCS <- exonTestability(RCS, cutoff=5)
> geneTestable <- geneTestability(RCS)
> RCS <- subsetByGenes(RCS, unique(geneID(RCS))[ geneTestable ])
> # get gene IDs, which will be used in initialization of gene set
> geneIDs <- unique(geneID(RCS))
> # calculate DS NB statistics
```

```

> RCS <- estiExonNBstat(RCS)
> RCS <- estiGeneNBstat(RCS)
> # calculate DS NB statistics on the permutation data sets
> permuteMat <- genpermuteMat(RCS, times=perm.times)
> RCS <- DSpermute4GSEA(RCS, permuteMat)

```

Step 2: DE analysis

```

> # get gene read counts
> geneCounts <- getGeneCount(RCS)
> # calculate DE NB statistics
> label <- label(RCS)
> dds <- runDESeq(geneCounts, label)
> DEGres <- DENBStat4GSEA(dds)
> # calculate DE NB statistics on the permutation data sets
> DEpermNBstat <- DENBStatPermut4GSEA(dds, permuteMat) # permutation

```

Step 3: score integration

```

> # DE score normalization
> DEScore.normFac <- normFactor(DEpermNBstat)
> DEScore <- scoreNormalization(DEGres$NBstat, DEScore.normFac)
> DEScore.perm <- scoreNormalization(DEpermNBstat, DEScore.normFac)
> # DS score normalization
> DSscore.normFac <- normFactor(RCS@permute_NBstat_gene)
> DSscore <- scoreNormalization(RCS@featureData_gene$NBstat, DSscore.normFac)
> DSscore.perm <- scoreNormalization(RCS@permute_NBstat_gene, DSscore.normFac)
> # score integration
> gene.score <- geneScore(DEscore, DSscore, DEweight=0.5)
> gene.score.perm <- genePermuteScore(DEscore.perm, DSscore.perm, DEweight=0.5)
> # visilization of scores
> # NOT run in the example; users to uncomment the following 6 lines to run
> #plotGeneScore(DEscore, DEScore.perm, pdf=paste(output.prefix, ".DEScore.pdf", sep=""),
> #             main="Expression")
> #plotGeneScore(DSscore, DSscore.perm, pdf=paste(output.prefix, ".DSScore.pdf", sep=""),
> #             main="Splicing")
> #plotGeneScore(gene.score, gene.score.perm,
> #             pdf=paste(output.prefix, ".GeneScore.pdf", sep=""))

```

Step 4: main GSEA

```

> # load gene set data
> gene.set <- loadGenesets(geneset.file, geneIDs, geneID.type="ensembl",
+                       genesetsize.min = 5, genesetsize.max = 1000)
> # enrichment analysis
> gene.set <- GSEnrichAnalyze(gene.set, gene.score, gene.score.perm, weighted.type=1)
> # format enrichment analysis results
> GSEAres <- GSEAresultTable(gene.set, TRUE)
> # output results
> # NOT run in the example; users to uncomment the following 4 lines to run
> #write.table(GSEAres, paste(output.prefix, ".GSEA.result.txt", sep=""),
> #          quote=FALSE, sep="\t", row.names=FALSE)
> #plotES(gene.set, pdf=paste(output.prefix, ".GSEA.ES.pdf", sep=""))
> #plotSig(gene.set, pdf=paste(output.prefix, ".GSEA.FDR.pdf", sep=""))

```


For gene sets used in Step 4, while we recommend users directly download and use those already well-formatted gene sets from MSigDB (<http://www.broadinstitute.org/gsea/msigdb/index.jsp>), users can also feed whatever gene sets to SeqGSEA as long as they are in the GMT format. Please refer to the following URL for details: http://www.broadinstitute.org/cancer/software/gsea/wiki/index.php/Data_formats.

6.3 Exemplified pipeline for DE-only analysis

For the demanding of DE-only analysis, such as for organisms without much alternative splicing annotated, here we show an exemplified pipeline for such analysis. It includes 4 steps as follows.

Step 0: Initialization. (Users should change values in this part accordingly.)

```
> rm(list=ls())
> # input count data files
> data.dir <- system.file("extdata", package="SeqGSEA", mustWork=TRUE)
> count.file <- paste(data.dir, "geneCounts.txt", sep="/")
> # gene set file
> geneset.file <- system.file("extdata", "gs_symb.txt",
+                             package="SeqGSEA", mustWork=TRUE)
> # output file prefix
> output.prefix <- "SeqGSEA.test"
> # setup parallel backend
> library(doParallel)
> cl <- makeCluster(2) # specify 2 cores to be used in computing
> registerDoParallel(cl) # parallel backend registration
> # setup permutation times
> perm.times <- 20 # change the number to >= 1000 in your analysis
```

Step 1: DE analysis

```
> # load gene read count data
> geneCounts <- read.table(count.file)
> # specify the labels of each sample
> label <- as.factor(c(rep(1,10), rep(0,10)))
> # calculate DE NB statistics
> dds <- runDESeq(geneCounts, label)
> DEGres <- DENBStat4GSEA(dds)
> # calculate DE NB statistics on the permutation data sets
> permuteMat <- genpermuteMat(label, times=perm.times)
> DEpermNBstat <- DENBStatPermut4GSEA(dds, permuteMat) # permutation
```

Step 2: score normalization

```
> # DE score normalization
> DEScore.normFac <- normFactor(DEpermNBstat)
> DEScore <- scoreNormalization(DEGres$NBstat, DEScore.normFac)
> DEScore.perm <- scoreNormalization(DEpermNBstat, DEScore.normFac)
> # score integration - DSscore can be null
> gene.score <- geneScore(DEscore, DEweight=1)
> gene.score.perm <- genePermuteScore(DEscore.perm, DEweight=1) # visualization of scores
> # NOT run in the example; users to uncomment the following 6 lines to run
> #plotGeneScore(DEscore, DEScore.perm, pdf=paste(output.prefix, ".DEscore.pdf", sep=""),
> #              main="Expression")
> #plotGeneScore(gene.score, gene.score.perm,
> #              pdf=paste(output.prefix, ".GeneScore.pdf", sep=""))
```

Step 3: main GSEA

```
> # load gene set data
> geneIDs <- rownames(geneCounts)
> gene.set <- loadGenesets(geneset.file, geneIDs, geneID.type="ensembl",
+                         genesetsize.min = 5, genesetsize.max = 1000)
> # enrichment analysis
> gene.set <- GSEnrichAnalyze(gene.set, gene.score, gene.score.perm, weighted.type=1)
> # format enrichment analysis results
> GSEAs <- GSEAsresultTable(gene.set, TRUE)
> # output results
> # NOT run in the example; users to uncomment the following 4 lines to run
> #write.table(GSEAs, paste(output.prefix, ".GSEA.result.txt", sep=""),
> #           quote=FALSE, sep="\t", row.names=FALSE)
> #plotES(gene.set, pdf=paste(output.prefix, ".GSEA.ES.pdf", sep=""))
> #plotSig(gene.set, pdf=paste(output.prefix, ".GSEA.FDR.pdf", sep=""))
```

6.4 One-step SeqGSEA analysis

While users can choose to run SeqGSEA step by step in a well-controlled manner (see above), the one-step SeqGSEA analysis with an all-in `runSeqGSEA` function enables users to run SeqGSEA in the easiest way. With the `runSeqGSEA` function, users can also test multiple weights for integrating DE and DS scores. DE-only analysis starting with exon read counts is also supported in the all-in function.

Follow the example below to start your first SeqGSEA analysis now!

```
> ### Initialization ###
> # input file location and pattern
> data.dir <- system.file("extdata", package="SeqGSEA", mustWork=TRUE)
> case.pattern <- "~SC" # file name starting with "SC"
> ctrl.pattern <- "~SN" # file name starting with "SN"
> # gene set file and type
> geneset.file <- system.file("extdata", "gs_symb.txt",
+                             package="SeqGSEA", mustWork=TRUE)
> geneID.type <- "ensembl"
> # output file prefix
> output.prefix <- "SeqGSEA.example"
> # analysis parameters
> nCores <- 8
> perm.times <- 1000 # >= 1000 recommended
> DEonly <- FALSE
> DEweight <- c(0.2, 0.5, 0.8) # a vector for different weights
> integrationMethod <- "linear"
>
> ### one step SeqGSEA running ###
> # NOT run in the example; uncomment the following 4 lines to run
> # CAUTION: running the following lines will generate lots of files in your working dir
> #runSeqGSEA(data.dir=data.dir, case.pattern=case.pattern, ctrl.pattern=ctrl.pattern,
> #          geneset.file=geneset.file, geneID.type=geneID.type, output.prefix=output.prefix,
> #          nCores=nCores, perm.times=perm.times, integrationMethod=integrationMethod,
> #          DEonly=DEonly, DEweight=DEweight)
```

7 Session information

```
> sessionInfo()
```

```
R Under development (unstable) (2024-10-21 r87258)
```

```
Platform: x86_64-pc-linux-gnu
```

```
Running under: Ubuntu 24.04.1 LTS
```

```
Matrix products: default
```

```
BLAS: /home/biocbuild/bbs-3.21-bioc/R/lib/libRblas.so
```

```
LAPACK: /usr/lib/x86_64-linux-gnu/lapack/liblapack.so.3.12.0
```

```
locale:
```

```
[1] LC_CTYPE=en_US.UTF-8      LC_NUMERIC=C
[3] LC_TIME=en_GB             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
```

```
time zone: America/New_York
```

```
tzcode source: system (glibc)
```

```
attached base packages:
```

```
[1] stats4      parallel  stats      graphics  grDevices  utils      datasets
[8] methods     base
```

```
other attached packages:
```

```
[1] SeqGSEA_1.47.0           DESeq2_1.47.0
[3] SummarizedExperiment_1.37.0 MatrixGenerics_1.19.0
[5] matrixStats_1.4.1        GenomicRanges_1.59.0
[7] GenomeInfoDb_1.43.0      IRanges_2.41.0
[9] S4Vectors_0.45.0         doParallel_1.0.17
[11] iterators_1.0.14         foreach_1.5.2
[13] Biobase_2.67.0           BiocGenerics_0.53.0
```

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

```
[1] KEGGREST_1.47.0          gtable_0.3.6             httr2_1.0.5
[4] ggplot2_3.5.1            lattice_0.22-6           vctrs_0.6.5
[7] tools_4.5.0              generics_0.1.3          curl_5.2.3
[10] tibble_3.2.1             fansi_1.0.6              AnnotationDbi_1.69.0
[13] RSQLite_2.3.7            blob_1.2.4               pkgconfig_2.0.3
[16] Matrix_1.7-1             dbplyr_2.5.0            lifecycle_1.0.4
[19] GenomeInfoDbData_1.2.13 stringr_1.5.1            compiler_4.5.0
[22] progress_1.2.3           Biostrings_2.75.0       munsell_0.5.1
[25] codetools_0.2-20        pillar_1.9.0            crayon_1.5.3
[28] BiocParallel_1.41.0     DelayedArray_0.33.0     cachem_1.1.0
[31] abind_1.4-8              digest_0.6.37           tidyselect_1.2.1
[34] locfit_1.5-9.10         stringi_1.8.4           purrr_1.0.2
[37] dplyr_1.1.4             biomaRt_2.63.0         fastmap_1.2.0
[40] grid_4.5.0              colorspace_2.1-1        cli_3.6.3
[43] SparseArray_1.7.0       magrittr_2.0.3          S4Arrays_1.7.0
[46] utf8_1.2.4              withr_3.0.2             filelock_1.0.3
```

[49]	rappdirs_0.3.3	prettyunits_1.2.0	scales_1.3.0
[52]	UCSC.utils_1.3.0	bit64_4.5.2	XVector_0.47.0
[55]	httr_1.4.7	bit_4.5.0	hms_1.1.3
[58]	png_0.1-8	memoise_2.0.1	BiocFileCache_2.15.0
[61]	rlang_1.1.4	Rcpp_1.0.13	glue_1.8.0
[64]	DBI_1.2.3	xml2_1.3.6	jsonlite_1.8.9
[67]	R6_2.5.1	zlibbioc_1.53.0	

Cleanup

This is a cleanup step for the vignette on Windows; typically not needed for users.

```
> allCon <- showConnections()
> socketCon <- as.integer(rownames(allCon)[allCon[, "class"] == "sockconn"])
> sapply(socketCon, function(ii) close.connection(getConnection(ii)) )
```

References

- Anders, S. and Huber, W. (2010). Differential expression analysis for sequence count data. *Genome Biology*, 11:R106.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B*, 57(1):289–300.
- Kannan, K., Wang, L., Wang, J., Ittmann, M. M., Li, W., and Yen, L. (2011). Recurrent chimeric rnas enriched in human prostate cancer identified by deep sequencing. *Proc Natl Acad Sci U S A*, 108(22):9172–7.
- Subramanian, A., Tamayo, P., Mootha, V. K., Mukherjee, S., Ebert, B. L., Gillette, M. A., Paulovich, A., Pomeroy, S. L., Golub, T. R., Lander, E. S., and Mesirov, J. P. (2005). Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci U S A*, 102(43):15545–50.
- Wang, W., Qin, Z., Feng, Z., Wang, X., and Zhang, X. (2013). Identifying differentially spliced genes from two groups of rna-seq samples. *Gene*, 518(1):164–170.
- Wang, X. and Cairns, M. (2013). Gene set enrichment analysis of RNA-Seq data: integrating differential expression and splicing. *BMC Bioinformatics*, 14(Suppl 5):S16.
- Wang, X. and Cairns, M. (2014). SeqGSEA: a Bioconductor package for gene set enrichment analysis of RNA-Seq data integrating differential expression and splicing. *Bioinformatics*, 30(12):1777–9.