

metagenomeSeq: Statistical analysis for sparse high-throughput sequencing

Joseph Nathaniel Paulson

Applied Mathematics & Statistics, and Scientific Computation
Center for Bioinformatics and Computational Biology
University of Maryland, College Park

`jpaulson@umiacs.umd.edu`

Modified: October 6, 2014. Compiled: December 18, 2014

Contents

1	Introduction	3
2	Data preparation	4
2.1	Example datasets	4
2.2	Loading count data	5
2.3	Loading taxonomy	6
2.4	Loading metadata	6
2.5	Creating a MRexperiment object	6
2.6	Useful commands	7
3	Normalization	10
3.1	Calculating normalization factors	10
3.2	Exporting data	10
4	Statistical testing	12
4.1	Zero-inflated Gaussian mixture model	12
4.1.1	Example using fitZig for differential abundance testing	12
4.1.2	Multiple groups	14
4.1.3	Exporting fits	15
4.2	Time series analysis	16
4.2.1	Time series plotting	18
4.3	Permutation test	19
4.4	Presence-absence testing	20
4.5	Discovery odds ratio testing	20
4.6	Feature correlations	21
5	Aggregating features	22

6	Visualization of features	23
6.1	Interactive Display	23
6.2	Structural overview	23
6.3	Feature specific	25
7	Summary	28
7.1	Citing metagenomeSeq	28
7.2	Session Info	28
8	Appendix	30
8.1	Appendix A: MRexperiment internals	30
8.2	Appendix B: Mathematical model	30
8.3	Appendix C: Calculating the proper percentile	30

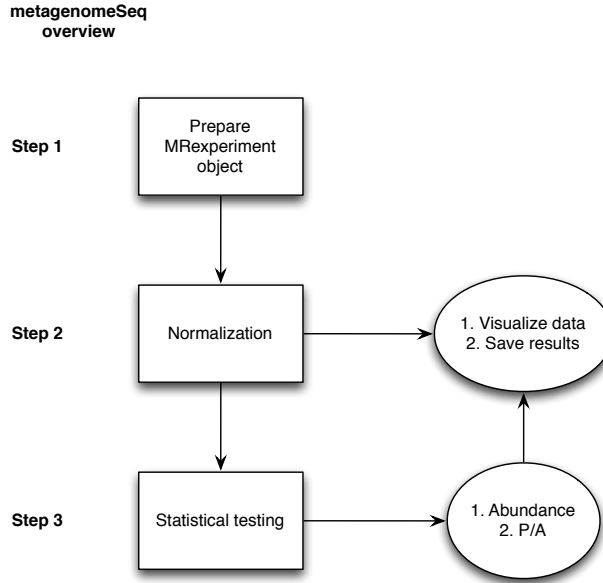


Figure 1: General overview. `metagenomeSeq` requires the user to convert their data into MR-experiment objects. Using those MRexperiment objects, one can normalize their data, run statistical tests (abundance or presence-absence), and visualize or save results.

1 Introduction

This is a vignette for pieces of an association study pipeline. For a full list of functions available in the package: `help(package=metagenomeSeq)`. For more information about a particular function call: `?function`.

Metagenomics is the study of genetic material targeted directly from an environmental community. Originally focused on exploratory and validation projects, these studies now focus on understanding the differences in microbial communities caused by phenotypic differences. Analyzing high-throughput sequencing data has been a challenge to researchers due to the unique biological and technological biases that are present in marker-gene survey data.

We present a R package, `metagenomeSeq`, that implements methods developed to account for previously unaddressed biases specific to high-throughput sequencing microbial marker-gene survey data. Our method implements a novel normalization technique and method to account for sparsity due to undersampling. Other methods include White *et al.*'s Metastats and Segata *et al.*'s LEfSe. The first is a non-parametric permutation test on *t*-statistics and the second is a non-parametric Kruskal-Wallis test followed by subsequent wilcox rank-sum tests on subgroups to guard against positive discoveries of differential abundance driven by potential confounders - neither address normalization nor sparsity.

This vignette describes the basic protocol when using `metagenomeSeq`. A normalization method able to control for biases in measurements across taxonomic features and a mixture model that implements a zero-inflated Gaussian distribution to account for varying depths of coverage are implemented. Using a linear model methodology, it is easy to include confounding sources of variability and interpret results. Additionally, visualization functions are provided to examine discoveries.

The software was designed to determine features (be it Operational Taxonomic Unit (OTU), species, etc.) that are differentially abundant between two or more groups of multiple samples. The software was also designed to address the effects of both normalization and undersampling of microbial communities on disease association detection and testing of feature correlations.

2 Data preparation

Microbial marker gene sequence data is preprocessed and counts are algorithmically defined from project-specific sequence data by clustering reads according to read similarity. Given m features and n samples, the elements in a count matrix \mathbf{C} (m, n), c_{ij} , are the number of reads annotated for a particular feature i (whether it be OTU, species, genus, etc.) in sample j .

$$\begin{matrix} & \text{sample}_1 & \text{sample}_2 & \dots & \text{sample}_n \\ \text{feature}_1 & \left(\begin{matrix} c_{11} & c_{12} & \dots & c_{1n} \\ \text{feature}_2 & c_{21} & c_{22} & \dots & c_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \text{feature}_m & c_{m1} & c_{m2} & \dots & c_{mn} \end{matrix} \right) \end{matrix}$$

Count data should be stored in a delimited (tab by default) file with sample names along the first row and feature names along the first column.

Data is prepared and formatted as a MRexperiment object. For an overview of the internal structure please see Appendix A.

2.1 Example datasets

There are two datasets included as examples in the metagenomeSeq package. Data needs to be in a MRexperiment object format to normalize, run statistical tests, and visualize. As an example, throughout the vignette we'll use the following datasets. To understand a function's usage or included data simply enter ?functionName.

```
library(metagenomeSeq)
```

1. Human lung microbiome [1]: The lung microbiome consists of respiratory flora sampled from six healthy individuals. Three healthy nonsmokers and three healthy smokers. The upper lung tracts were sampled by oral wash and oro-/nasopharyngeal swabs. Samples were taken using two bronchoscopes, serial bronchoalveolar lavage and lower airway protected brushes.

```
data(lungData)
lungData

## MRexperiment (storageMode: environment)
## assayData: 51891 features, 78 samples
##   element names: counts
## protocolData: none
## phenoData
##   sampleNames: CHK_6467_E3B11_BRONCH2_PREWASH_V1V2
##                 CHK_6467_E3B11_OW_V1V2 ... CHK_6467_E3B09_BAL_A_V1V2
##                 (78 total)
##   varLabels: SampleType SiteSampled SmokingStatus
##   varMetadata: labelDescription
## featureData
##   featureNames: 1 2 ... 51891 (51891 total)
##   fvarLabels: taxa
```

```
## fvarMetadata: labelDescription
## experimentData: use 'experimentData(object)'
## PubMedIds: 21680950
## Annotation:
```

2. Humanized gnotobiotic mouse gut [2]: Twelve germ-free adult male C57BL/6J mice were fed a low-fat, plant polysaccharide-rich diet. Each mouse was gavaged with healthy adult human fecal material. Following the fecal transplant, mice remained on the low-fat, plant polysaccharide-rich diet for four weeks, following which a subset of 6 were switched to a high-fat and high-sugar diet for eight weeks. Fecal samples for each mouse went through PCR amplification of the bacterial 16S rRNA gene V2 region weekly. Details of experimental protocols and further details of the data can be found in Turnbaugh et. al. Sequences and further information can be found at: http://gordonlab.wustl.edu/TurnbaughSE_10_09/STM_2009.html

```
data(mouseData)
mouseData

## MRExperiment (storageMode: environment)
## assayData: 10172 features, 139 samples
## element names: counts
## protocolData: none
## phenoData
## sampleNames: PM1:20080107 PM1:20080108 ... PM9:20080303
## (139 total)
## varLabels: mouseID date ... status (5 total)
## varMetadata: labelDescription
## featureData
## featureNames: Prevotellaceae:1 Lachnospiraceae:1 ...
## Parabacteroides:956 (10172 total)
## fvarLabels: taxa class
## fvarMetadata: labelDescription
## experimentData: use 'experimentData(object)'
## Annotation:
```

2.2 Loading count data

Following preprocessing and annotation of sequencing data metagenomeSeq requires a count matrix with features along rows and samples along the columns. metagenomeSeq includes functions for loading delimited files of counts `load_meta` and phenodata `load_phenoData`.

As an example, a portion of the lung microbiome [1] OTU matrix is provided in metagenomeSeq's library "extdata" folder. The OTU matrix is stored as a tab delimited file. `load_meta` loads the taxa and counts into a list.

```
dataDirectory <- system.file("extdata", package = "metagenomeSeq")
lung = load_meta(file.path(dataDirectory, "CHK_NAME.otus.count.csv"))
dim(lung$counts)

## [1] 1000 78
```

2.3 Loading taxonomy

Next we want to load the annotated taxonomy. Check to make sure that your taxa annotations and OTUs are in the same order as your matrix rows.

```
taxa = read.delim(file.path(dataDirectory, "CHK_otus.taxonomy.csv"),
  stringsAsFactors = FALSE)
```

As our OTUs appear to be in order with the count matrix we loaded earlier, the next step is to load phenodata.

Warning: features need to have the same names as the rows of the count matrix when we create the MRexperiment object for provenance purposes.

2.4 Loading metadata

Phenotype data can be optionally loaded into R with `load_phenoData`. This function loads the data as a list.

```
clin = load_phenoData(file.path(dataDirectory, "CHK_clinical.csv"),
  tran = TRUE)
ord = match(colnames(lung$counts), rownames(clin))
clin = clin[ord, ]
head(clin[1:2, ])

##                               SampleType
## CHK_6467_E3B11_BRONCH2_PREWASH_V1V2 Bronch2.PreWash
## CHK_6467_E3B11_OW_V1V2                               OW
##                               SiteSampled
## CHK_6467_E3B11_BRONCH2_PREWASH_V1V2 Bronchoscope.Channel
## CHK_6467_E3B11_OW_V1V2                               OralCavity
##                               SmokingStatus
## CHK_6467_E3B11_BRONCH2_PREWASH_V1V2           Smoker
## CHK_6467_E3B11_OW_V1V2                       Smoker
```

Warning: phenotypes must have the same names as the columns on the count matrix when we create the MRexperiment object for provenance purposes.

2.5 Creating a MRexperiment object

Function `newMRexperiment` takes a count matrix, `phenoData` (annotated data frame), and `featureData` (annotated data frame) as input. Biobase provides functions to create annotated data frames. Library sizes (depths of coverage) and normalization factors are also optional inputs.

```
phenotypeData = AnnotatedDataFrame(clin)
phenotypeData

## An object of class 'AnnotatedDataFrame'
##   rowNames: CHK_6467_E3B11_BRONCH2_PREWASH_V1V2
##             CHK_6467_E3B11_OW_V1V2 ... CHK_6467_E3B09_BAL_A_V1V2
##             (78 total)
##   varLabels: SampleType SiteSampled SmokingStatus
##   varMetadata: labelDescription
```

A feature annotated data frame. In this example it is simply the OTU numbers, but it can as easily be the annotated taxonomy at multiple levels.

```
OTUdata = AnnotatedDataFrame(taxa)
OTUdata

## An object of class 'AnnotatedDataFrame'
##   rowNames: 1 2 ... 1000 (1000 total)
##   varLabels: OTU Taxonomy ... strain (10 total)
##   varMetadata: labelDescription

obj = newMRexperiment(lung$counts, phenoData=phenotypeData, featureData=OTUdata)
# Links to a paper providing further details can be included optionally.
# experimentData(obj) = annotate::pmid2MIAME("21680950")
obj

## MRexperiment (storageMode: environment)
## assayData: 1000 features, 78 samples
##   element names: counts
## protocolData: none
## phenoData
##   sampleNames: CHK_6467_E3B11_BRONCH2_PREWASH_V1V2
##                 CHK_6467_E3B11_OW_V1V2 ... CHK_6467_E3B09_BAL_A_V1V2
##                 (78 total)
##   varLabels: SampleType SiteSampled SmokingStatus
##   varMetadata: labelDescription
## featureData
##   featureNames: 1 2 ... 1000 (1000 total)
##   fvarLabels: OTU Taxonomy ... strain (10 total)
##   fvarMetadata: labelDescription
## experimentData: use 'experimentData(object)'
## Annotation:
```

Alternatively, you can load in Biome-format data (outputs of QIIME and Mothur) using the `loadbiom` function. If a biom class object is already loaded into R it can be converted to a MRexperiment-class object using the `biom2MRexperiment` function.

2.6 Useful commands

Phenotype information can be accessed with the `phenoData` and `pData` methods:

```
phenoData(obj)

## An object of class 'AnnotatedDataFrame'
##   sampleNames: CHK_6467_E3B11_BRONCH2_PREWASH_V1V2
##                 CHK_6467_E3B11_OW_V1V2 ... CHK_6467_E3B09_BAL_A_V1V2
##                 (78 total)
##   varLabels: SampleType SiteSampled SmokingStatus
##   varMetadata: labelDescription

head(pData(obj), 3)
```

```
##                                     SampleType
## CHK_6467_E3B11_BRONCH2_PREWASH_V1V2 Bronch2.PreWash
## CHK_6467_E3B11_OW_V1V2                               OW
## CHK_6467_E3B08_OW_V1V2                               OW
##                                     SiteSampled
## CHK_6467_E3B11_BRONCH2_PREWASH_V1V2 Bronchoscope.Channel
## CHK_6467_E3B11_OW_V1V2                               OralCavity
## CHK_6467_E3B08_OW_V1V2                               OralCavity
##                                     SmokingStatus
## CHK_6467_E3B11_BRONCH2_PREWASH_V1V2 Smoker
## CHK_6467_E3B11_OW_V1V2 Smoker
## CHK_6467_E3B08_OW_V1V2 NonSmoker
```

Feature information can be accessed with the `featureData` and `fData` methods:

```
featureData(obj)

## An object of class 'AnnotatedDataFrame'
##   featureNames: 1 2 ... 1000 (1000 total)
##   varLabels: OTU Taxonomy ... strain (10 total)
##   varMetadata: labelDescription

head(fData(obj)[, -c(2, 10)], 3)

##   OTU superkingdom      phylum      class
## 1   1      Bacteria Proteobacteria Epsilonproteobacteria
## 2   2      <NA>          <NA>          <NA>
## 3   3      Bacteria Actinobacteria Actinobacteria (class)
##                                     order      family      genus
## 1 Campylobacterales Campylobacteraceae Campylobacter
## 2      <NA>          <NA>          <NA>
## 3 Actinomycetales Actinomycetaceae Actinomyces
##                                     species
## 1      Campylobacter rectus
## 2      <NA>
## 3 Actinomyces radicidentis
```

The raw or normalized counts matrix can be accessed with the `MRcounts` function:

```
head(MRcounts(obj[, 1:2]))

##   CHK_6467_E3B11_BRONCH2_PREWASH_V1V2 CHK_6467_E3B11_OW_V1V2
## 1                                     0                          0
## 2                                     0                          0
## 3                                     0                          0
## 4                                     0                          0
## 5                                     0                          0
## 6                                     0                          0
```

A `MRexperiment`-class object can be easily subsetted, for example:


```

featuresToKeep = which(rowSums(obj) >= 100)
samplesToKeep = which(pData(obj)$SmokingStatus == "Smoker")
obj_smokers = obj[featuresToKeep, samplesToKeep]
obj_smokers

## MRExperiment (storageMode: environment)
## assayData: 1 features, 33 samples
##   element names: counts
## protocolData: none
## phenoData
##   sampleNames: CHK_6467_E3B11_BRONCH2_PREWASH_V1V2
##                 CHK_6467_E3B11_OW_V1V2 ... CHK_6467_E3B09_BAL_A_V1V2
##                 (33 total)
##   varLabels: SampleType SiteSampled SmokingStatus
##   varMetadata: labelDescription
## featureData
##   featureNames: 570
##   fvarLabels: OTU Taxonomy ... strain (10 total)
##   fvarMetadata: labelDescription
## experimentData: use 'experimentData(object)'
## Annotation:

head(pData(obj_smokers), 3)

##                                     SampleType
## CHK_6467_E3B11_BRONCH2_PREWASH_V1V2 Bronch2.PreWash
## CHK_6467_E3B11_OW_V1V2                               OW
## CHK_6467_E3B11_BAL_A_V1V2                           BAL.A
##                                     SiteSampled
## CHK_6467_E3B11_BRONCH2_PREWASH_V1V2 Bronchoscope.Channel
## CHK_6467_E3B11_OW_V1V2                               OralCavity
## CHK_6467_E3B11_BAL_A_V1V2                           Lung
##                                     SmokingStatus
## CHK_6467_E3B11_BRONCH2_PREWASH_V1V2      Smoker
## CHK_6467_E3B11_OW_V1V2                   Smoker
## CHK_6467_E3B11_BAL_A_V1V2                Smoker

```

3 Normalization

Normalization is required due to varying depths of coverage across samples. `cumNorm` is a normalization method that calculates scaling factors equal to the sum of counts up to a particular quantile.

Denote the l th quantile of sample j as q_j^l , that is, in sample j there are l taxonomic features with counts smaller than q_j^l . For $l = \lfloor .95m \rfloor$ then q_j^l corresponds to the 95th percentile of the count distribution for sample j .

Denote $s_j^l = \sum_{(i|c_{ij} \leq q_j^l)} c_{ij}$ as the sum of counts for sample j up to the l th quantile. Our normalization chooses a value $\hat{l} \leq m$ to define a normalization scaling factor for each sample to produce normalized counts $\tilde{c}_{ij} = \frac{c_{ij}}{s_j^{\hat{l}}} N$ where N is an appropriately chosen normalization constant. See Appendix C for more information on how our method calculates the proper percentile.

These normalization factors are stored in the experiment summary slot. Functions to determine the proper percentile `cumNormStat`, save normalized counts `exportMat`, or save various sample statistics `exportStats` are also provided. Normalized counts can be called easily by `cumNormMat` (`MRExperimentObject`) or `MRcounts` (`MRExperimentObject`, `norm=TRUE`, `log=FALSE`)

3.1 Calculating normalization factors

After defining a `MRExperiment` object, the first step is to calculate the proper percentile by which to normalize counts. There are several options in calculating and visualizing the relative differences in the reference. Figure 3 is an example from the lung dataset.

```
data(lungData)
p = cumNormStatFast(lungData)
```

To calculate the scaling factors we simply run `cumNorm`

```
lungData = cumNorm(lungData, p = p)
```

The user can alternatively choose different percentiles for the normalization scheme by specifying p .

There are other functions, including `normFactors`, `cumNormMat`, that return the normalization factors or a normalized matrix for a specified percentile. To see a full list of functions please refer to the manual and help pages.

3.2 Exporting data

To export normalized count matrices:

```
mat = MRcounts(lungData, norm = TRUE, log = TRUE)[1:5, 1:5]
exportMat(mat, file = file.path(dataDirectory, "tmp.tsv"))
```

To save sample statistics (sample scaling factor, quantile value, number of identified features and library size):

```
exportStats(lungData[, 1:5], file = file.path(dataDirectory, "tmp.tsv"))
head(read.csv(file = file.path(dataDirectory, "tmp.tsv"), sep = "\t"))
```

##	Subject	Scaling.factor
## 1	CHK_6467_E3B11_BRONCH2_PREWASH_V1V2	67
## 2	CHK_6467_E3B11_OW_V1V2	2475
## 3	CHK_6467_E3B08_OW_V1V2	2198
## 4	CHK_6467_E3B07_BAL_A_V1V2	836
## 5	CHK_6467_E3B11_BAL_A_V1V2	1008

##	Quantile.value	Number.of.identified.features	Library.size
## 1	2	60	271
## 2	1	3299	7863
## 3	1	2994	8360
## 4	1	1188	5249
## 5	1	1098	3383




Figure 2: The number of unique features is plotted against depth of coverage for samples from the Human Microbiome Project [3]. Including the depth of coverage and the interaction of body site and sequencing site we are able to achieve an adjusted R^2 of .94. The zero-inflated Gaussian mixture was developed to account for missing features.

4 Statistical testing

Now that we have taken care of normalization we can address the effects of under sampling on the detecting differentially abundant features (OTUs, genes, etc).

4.1 Zero-inflated Gaussian mixture model

The depth of coverage in a sample is directly related to how many features are detected in a sample motivating our zero-inflated Gaussian (ZIG) mixture model. Figure 2 is representative of the linear relationship between depth of coverage and OTU identification ubiquitous in marker-gene survey datasets currently available. For a quick overview of the mathematical model see Appendix B.

Function `fitZig` performs a complex mathematical optimization routine to estimate probabilities that a zero for a particular feature in a sample is a technical zero or not. The function relies heavily on the `limma` package [4]. Design matrices can be created in R by using the `model.matrix` function and are inputs for `fitZig`.

For large survey studies it is often pertinent to include phenotype information or confounders into a design matrix when testing the association between the abundance of taxonomic features and a phenotype of interest (disease, for instance). Our linear model methodology can easily incorporate these confounding covariates in a straightforward manner. `fitZig` output includes weighted fits for each of the m features. Results can be filtered and saved using `MRcoefs` or `MRtable`.

4.1.1 Example using `fitZig` for differential abundance testing

Warning: The user should restrict significant features to those with a minimum number of positive samples. What this means is that one should not claim features are significant unless the effective number of samples is above a particular percentage. For example, fold-change




Figure 3: Relative difference for the median difference in counts from the reference.

estimates might be unreliable if an entire group does not have a positive count for the feature in question.

We recommend the user remove features based on the number of estimated effective samples, please see `calculateEffectiveSamples`. We recommend removing features with less than the average number of effective samples in all features. In essence, setting `eff = .5` when using `MRcoefs`, `MRfulltable`, or `MRtable`. To find features absent from a group the function `uniqueFeatures` provides a table of the feature ids, the number of positive features and reads for each group.

In our analysis of the lung microbiome data, we can remove features that are not present in many samples, controls, and calculate the normalization factors. The user needs to decide which metadata should be included in the linear model.

```
controls = grep("Extraction.Control", pData(lungData)$SampleType)
lungTrim = lungData[, -controls]
sparseFeatures = which(rowSums(MRcounts(lungTrim) > 0) < 10)
lungTrim = lungTrim[-sparseFeatures, ]
lungp = cumNormStat(lungTrim, pFlag = TRUE, main = "Trimmed lung data")

lungTrim = cumNorm(lungTrim, p = lungp)
```

After the user defines an appropriate model matrix for hypothesis testing there are optional inputs to `fitZig`, including settings determined by `zigControl`. We ask the user to review the help files for both `fitZig` and `zigControl`. For this example we include body site as covariates and want to test for the bacteria differentially abundant between smokers and non-smokers.

```
smokingStatus = pData(lungTrim)$SmokingStatus
bodySite = pData(lungTrim)$SampleType
normFactor = normFactors(lungTrim)
```

```

normFactor = log2(normFactor/median(normFactor) + 1)
mod = model.matrix(~smokingStatus + bodySite + normFactor)
settings = zigControl(maxit = 10, verbose = TRUE)
fit = fitZig(obj = lungTrim, mod = mod, useCSSoffset = FALSE, control = settings)

## it= 0, nll=88.42, log10(eps+1)=Inf, stillActive=1029
## it= 1, nll=93.56, log10(eps+1)=0.06, stillActive=261
## it= 2, nll=93.46, log10(eps+1)=0.05, stillActive=120
## it= 3, nll=93.80, log10(eps+1)=0.05, stillActive=22
## it= 4, nll=93.94, log10(eps+1)=0.03, stillActive=3
## it= 5, nll=93.93, log10(eps+1)=0.00, stillActive=1
## it= 6, nll=93.90, log10(eps+1)=0.00, stillActive=1
## it= 7, nll=93.87, log10(eps+1)=0.00, stillActive=1
## it= 8, nll=93.86, log10(eps+1)=0.00, stillActive=1
## it= 9, nll=93.85, log10(eps+1)=0.00, stillActive=1

# The default, useCSSoffset = TRUE, automatically includes the CSS
# scaling normalization factor.

```

The result, `fit`, is a list providing detailed estimates of the fits including a limma fit in `fit$fit` and an eBayes statistical fit in `fit$eb`. This data can be analyzed like any limma fit and in this example, the column of the fitted coefficients represents the fold-change for our "smoker" vs. "nonsmoker" analysis.

Looking at the particular analysis just performed, there appears to be OTUs representing two *Prevotella*, two *Neisseria*, a *Porphyromonas* and a *Leptotrichia* that are differentially abundant. One should check that similarly annotated OTUs are not equally differentially abundant in controls.

Alternatively, the user can input a model with their own normalization factors including them directly in the model matrix and specifying the option `useCSSoffset = FALSE` in `fitZig`.

4.1.2 Multiple groups

Assuming there are multiple groups it is possible to make use of Limma's `topTable` functions for F-tests and contrast functions to compare multiple groups and covariates of interest. The output of `fitZig` includes a 'MLArrayLM' Limma object that can be called on by other functions. When running `fitZig` by default there is an additional covariate added to the design matrix. The fit and the ultimate design matrix are crucial for contrasts.

```

settings = zigControl(maxit = 1, verbose = FALSE)
mod = model.matrix(~bodySite)
colnames(mod) = levels(bodySite)
# fitting the ZIG model
res = fitZig(obj = lungTrim, mod = mod, control = settings)
# The output of fitZig contains a list of various useful items.
# hint: names(res). Probably the most useful is the limma
# 'MLArrayLM' object called fit.
zigFit = res$fit
finalMod = res$fit$design

contrast.matrix = makeContrasts(BAL.A - BAL.B, OW - PSB, levels = finalMod)
fit2 = contrasts.fit(zigFit, contrast.matrix)
fit2 = eBayes(fit2)

```

```

topTable(fit2)

##          BAL.A...BAL.B  OW...PSB   AveExpr      F      P.Value
## 6901      0.173441382  1.466113  0.2435881  13.853265  1.185023e-05
## 6291     -0.106957350  1.658829  0.4671470  13.729227  1.289286e-05
## 40291     0.068929259  1.700238  0.2195735  13.190131  1.865332e-05
## 37977    -0.379954612  2.174071  0.4526060  12.809340  2.428240e-05
## 18531     0.373187924  2.075648  0.7343081  12.744687  2.540042e-05
## 40329    -0.071459976  1.481582  0.2475735  10.890870  9.518411e-05
## 7343     -0.228590783  1.559465  0.3116465  10.726007  1.073649e-04
## 36117    -0.286658827  2.233996  0.4084024  10.500278  1.267123e-04
## 46525    -0.003944065  1.643531  0.2636173  10.301332  1.467495e-04
## 7114      0.237252527  1.032275  0.2044220   9.464994  2.742563e-04
##          adj.P.Val
## 6901    0.005227407
## 6291    0.005227407
## 40291   0.005227407
## 37977   0.005227407
## 18531   0.005227407
## 40329   0.015782637
## 7343    0.015782637
## 36117   0.016298371
## 46525   0.016778354
## 7114    0.026096051

# See help pages on decideTests, topTable, topTableF, vennDiagram,
# etc.

```

Further specific details can be found in section 9.3 and beyond of the Limma user guide. The take home message is that to make use of any Limma functions one needs to extract the final model matrix used: `resfitdesign` and the `MLArrayLM` Limma fit object: `res$fit`.

4.1.3 Exporting fits

Currently functions are being developed to wrap and output results more neatly, but `MRcoefs`, `MRtable`, `MRfulltable` can be used to view coefficient fits and related statistics and export the data with optional output values - see help files to learn how they differ. An important note is that the `by` variable controls which coefficients are of interest whereas `coef` determines the display.

To only consider features that are found in a large percentage of effectively positive (positive samples + the weight of zero counts included in the Gaussian mixture) use the `eff` option in the `MRtables`.

```

taxa = sapply(strsplit(as.character(fData(lungTrim)$taxa), split = ";"),
  function(i) {
    i[length(i)]
  })
head(MRcoefs(fit, taxa = taxa, coef = 2))

##          smokingStatusSmoker
## Neisseria polysaccharea      -4.031612

```

```
## Neisseria meningitidis -3.958899
## Prevotella intermedia -2.927686
## Porphyromonas sp. UQD 414 -2.675306
## Prevotella paludivivens 2.575672
## Leptotrichia sp. oral clone FP036 2.574172
## pvalues adjPvalues
## Neisseria polysaccharea 6.064532e-15 2.496161e-13
## Neisseria meningitidis 1.617066e-14 6.162817e-13
## Prevotella intermedia 1.869400e-12 5.343369e-11
## Porphyromonas sp. UQD 414 8.705169e-10 1.279660e-08
## Prevotella paludivivens 2.767250e-08 2.230904e-07
## Leptotrichia sp. oral clone FP036 1.382567e-37 6.091832e-35
```

4.2 Time series analysis

Implemented in the `fitTimeSeries` function is a method for calculating time intervals for which bacteria are differentially abundant. Fitting is performed using Smoothing Splines ANOVA (SS-ANOVA), as implemented in the `gss` package. Given observations at multiple time points for two groups the method calculates a function modeling the difference in abundance across all time. Using group membership permutations we estimate a null distribution of areas under the difference curve for the time intervals of interest and report significant intervals of time.

Use of the function for analyses should cite: "Finding regions of interest in high throughput genomics data using smoothing splines" Talukder H, Paulson JN, Bravo HC. (Submitted)

```
# The gnotobiotic mice come from a longitudinal study ideal for
# this type of analysis. We choose to perform our analysis at the
# class level and look for differentially abundant time intervals
# for 'Actinobacteria'. For demonstrations sake we perform only
# 10 permutations.
res = fitTimeSeries(obj = mouseData, lvl = "class", feature = "Actinobacteria",
  class = "status", id = "mouseID", time = "relativeTime", B = 10)

## Loading required package: gss

# We observe a time period of differential abundance for
# 'Actinobacteria'
res$timeIntervals

##      Interval start Interval end      Area p.value
## [1,]              9          50 90.70142        0

str(res)

## List of 5
## $ timeIntervals: num [1, 1:4] 9 50 90.7 0
## ..- attr(*, "dimnames")=List of 2
## .. ..$ : NULL
## .. ..$ : chr [1:4] "Interval start" "Interval end" "Area" "p.value"
## $ data      : 'data.frame': 139 obs. of 4 variables:
## ..$ abundance: num [1:139] 0 3.82 3.13 7.4 0 ...
## ..$ class    : Factor w/ 2 levels "0","1": 1 1 1 1 1 1 1 1 1 1 ...
```



```
## ..$ time      : num [1:139] 21 22 28 0 35 6 42 49 56 63 ...
## ..$ id        : Factor w/ 12 levels "PM1","PM10","PM11",...: 1 1 1 1 1 1 1 1
## $ fit         : 'data.frame': 78 obs. of 3 variables:
## ..$ fit       : num [1:78] 0.401 0.537 0.67 0.8 0.928 ...
## ..$ se        : num [1:78] 1.078 1.015 0.96 0.912 0.87 ...
## ..$ timePoints: num [1:78] 0 1 2 3 4 5 6 7 8 9 ...
## $ perm        : num [1:10, 1] -2.18 -36.1 -84.33 1.5 88.5 ...
## $ call        : language fitSSTimeSeries(obj = obj, feature = feature, clas
```

For example, to test every class in the mouse dataset:

```
classes = unique(fData(mouseData)[, "class"])

timeSeriesFits = lapply(classes, function(i) {
  fitTimeSeries(obj=mouseData,
    feature=i,
    class="status",
    id="mouseID",
    time="relativeTime",
    lvl='class',
    B=1) # This should be set to a higher value
})
names(timeSeriesFits) = classes

# Removing classes of bacteria without a potentially
# interesting time interval difference.
timeSeriesFits = sapply(timeSeriesFits, function(i) {i[[1]]}) [-c(2, 9, 14)]

# Correcting for multiple testing.
for(i in 1:length(timeSeriesFits)){
  pvalues = timeSeriesFits[[i]][,4]
  adjPvalues = p.adjust(pvalues, "bonferroni")
  timeSeriesFits[[i]] = cbind(timeSeriesFits[[i]], adjPvalues)
}

# Naming the various interesting time intervals.
for(i in 1:length(timeSeriesFits)){
  rownames(timeSeriesFits[[i]]) =
    paste(
      paste(names(timeSeriesFits)[i], " interval", sep=""),
      1:nrow(timeSeriesFits[[i]]), sep=":"
    )
}

# Merging into a table.
timeSeriesFits = do.call(rbind, timeSeriesFits)
head(timeSeriesFits)

##               Interval start Interval end          Area
## Bacteroidetes interval:1          18          20      6.431405
## Bacteroidetes interval:2          22          72 -108.605389
## Bacteroidetes interval:3          76          77  -1.445231
```

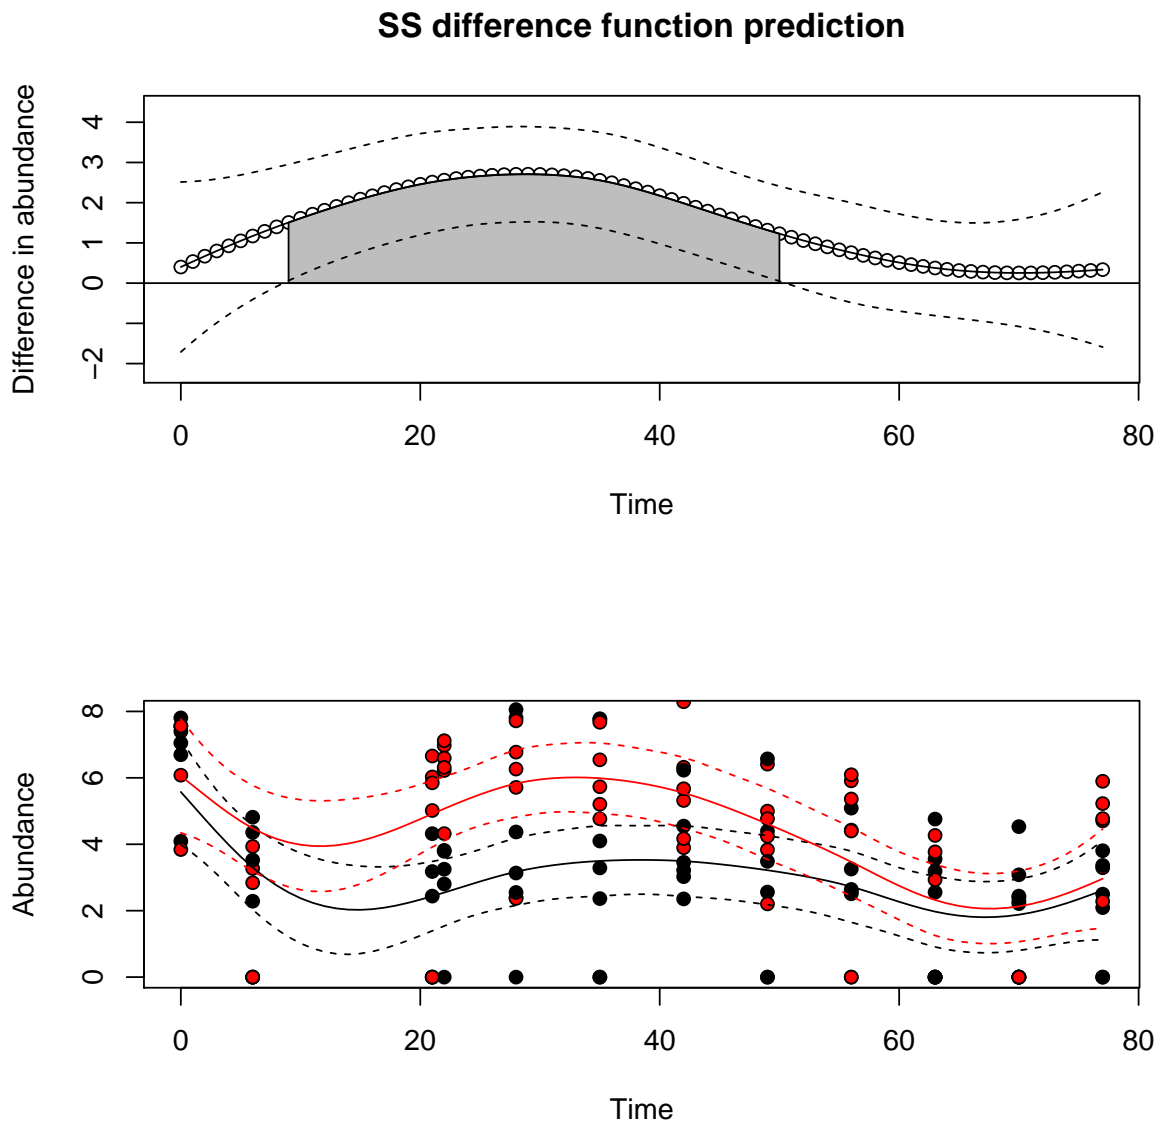
## Unknown interval:1	63	77	5.298119
## Unknown interval:2	0	22	-10.524746
## Bacilli interval:1	21	77	472.291720
##	p.value	adjPvalues	
## Bacteroidetes interval:1	0	0	
## Bacteroidetes interval:2	0	0	
## Bacteroidetes interval:3	0	0	
## Unknown interval:1	0	0	
## Unknown interval:2	0	0	
## Bacilli interval:1	0	0	

Please see the help page for `fitTimeSeries` for parameters. Note, only two groups can be compared to each other and the time parameter must be an actual value (currently no support for `posix`, etc.).

4.2.1 Time series plotting

We can plot the difference function fit and time intervals of differential abundance using `plotTimeSeries` and `plotClassTimeSeries` on the result. More plots will be updated in time.

```
par(mfrow = c(2, 1))
plotTimeSeries(res)
plotClassTimeSeries(res, pch = 21, bg = res$data$class, ylim = c(0,
8))
```



4.3 Permutation test

Included is a standard permutation test similar to what was used in Metastats. Below we show the fit for the same model as above using 10 permutations providing p-value resolution to the tenth. The `coef` parameter refers to the coefficient of interest to test. Below we first generate the list of significant features.

```

coeffOfInterest = 2
res = fitMeta(obj = lungTrim, mod = mod, useCSSoffset = FALSE, B = 10,
  coef = coeffOfInterest)

# extract p.values and adjust for multiple testing res$p are the
# p-values calculated through permutation
adjustedPvalues = p.adjust(res$p, method = "fdr")

# extract the absolute fold-change estimates
foldChange = abs(res$fit$coef[, coeffOfInterest])

```

```
# determine features still significant and order by the
sigList = which(adjustedPvalues <= 0.05)
sigList = sigList[order(foldChange[sigList])]

# view the top taxa associated with the coefficient of interest.
head(taxa[sigList])

## [1] "Prevotella salivae"          "OTU_21302"
## [3] "Prevotella sp. DJF_B116"    "Streptococcus dysgalactiae"
## [5] "Prevotella pallens"         "Prevotella sp. DJF_B116"
```

4.4 Presence-absence testing

The hypothesis for the implemented presence-absence test is that the proportion/odds of a given feature present is higher/lower among one group of individuals compared to another, and we want to test whether any difference in the proportions observed is significant. We use Fisher's exact test to create a 2x2 contingency table and calculate p-values, odd's ratios, and confidence intervals. `fitPA` calculates the presence-absence for each organism and returns a table of p-values, odd's ratios, and confidence intervals. The function will accept either a `MRExperiment` object or matrix. `MRfulltable` when sent a result of `fitZig` will also include the results of `fitPA`.

```
classes = pData(mouseData)$diet
res = fitPA(mouseData[1:5, ], cl = classes)
# Warning - the p-value is calculating 1 despite a high odd's
# ratio.
head(res)

##               oddsRatio      lower      upper      pvalues
## Prevotellaceae:1          Inf 0.01630496      Inf 1.0000000
## Lachnospiraceae:1          Inf 0.01630496      Inf 1.0000000
## Unclassified-Screened:1      Inf 0.01630496      Inf 1.0000000
## Clostridiales:1              0 0.00000000 24.77661 0.3884892
## Clostridiales:2          Inf 0.01630496      Inf 1.0000000
##               adjPvalues
## Prevotellaceae:1              1
## Lachnospiraceae:1              1
## Unclassified-Screened:1        1
## Clostridiales:1              1
## Clostridiales:2              1
```

4.5 Discovery odds ratio testing

The hypothesis for the implemented discovery test is that the proportion of observed counts for a feature of all counts are comparable between groups. We use Fisher's exact test to create a 2x2 contingency table and calculate p-values, odd's ratios, and confidence intervals. `fitDO` calculates the proportion of counts for each organism and returns a table of p-values, odd's ratios, and confidence intervals. The function will accept either a `MRExperiment` object or matrix.

```

classes = pData(mouseData)$diet
res = fitDO(mouseData[1:100, ], cl = classes, norm = FALSE, log = FALSE)
head(res)

##              oddsRatio      lower      upper      pvalues
## Prevotellaceae:1          Inf 0.01630496          Inf 1.0000000
## Lachnospiraceae:1          Inf 0.01630496          Inf 1.0000000
## Unclassified-Screened:1      Inf 0.01630496          Inf 1.0000000
## Clostridiales:1              0 0.00000000 24.77661 0.3884892
## Clostridiales:2          Inf 0.01630496          Inf 1.0000000
## Firmicutes:1              0 0.00000000 24.77661 0.3884892
##
##              adjPvalues
## Prevotellaceae:1      1.0000000
## Lachnospiraceae:1      1.0000000
## Unclassified-Screened:1 1.0000000
## Clostridiales:1      0.7470946
## Clostridiales:2      1.0000000
## Firmicutes:1          0.7470946

```

4.6 Feature correlations

To test the correlations of abundance features, or samples, in a pairwise fashion we have implemented `correlationTest` and `correctIndices`. The `correlationTest` function will calculate basic pearson, spearman, kendall correlation statistics for the rows of the input and report the associated p-values. If a vector of length `ncol(obj)` it will also calculate the correlation of each row with the associated vector.

```

cors = correlationTest(mouseData[55:60, ], norm = FALSE, log = FALSE)
head(cors)

##              correlation              p
## Clostridiales:11-Lachnospiraceae:35 -0.02205882 7.965979e-01
## Clostridiales:11-Coprobacillus:3     -0.01701180 8.424431e-01
## Clostridiales:11-Lactobacillales:3    -0.01264304 8.825644e-01
## Clostridiales:11-Enterococcaceae:3     0.57315130 1.663114e-13
## Clostridiales:11-Enterococcaceae:4    -0.01264304 8.825644e-01
## Lachnospiraceae:35-Coprobacillus:3     0.24572606 3.548360e-03

```

Caution: <http://www.ncbi.nlm.nih.gov/pubmed/23028285>

5 Aggregating features

Normalization is recommended at the OTU level. However, functions are in place to aggregate the count matrix (normalized or not), based on a particular user defined level. Using the feature-Data information in the MRexperiment object, calling `aggregateByTaxonomy` or `aggTax` on a MRexperiment object and declaring particular featureData column name (i.e. 'genus') will aggregate counts to the desired level with the `aggfun` function (default `colSums`). Possible `aggfun` alternatives include `colMeans` and `colMedians`.

```
# first featureData column as 'chr' vector
taxa = as.character(fData(mouseData)[, 1])

# 2nd item of each taxonomy if delimited by ';'
phylum = sapply(strsplit(taxa, split = ";"), function(i) {
  i[2]
})

obj = aggTax(MRcounts(mouseData), lvl = phylum, out = "matrix")
head(obj[1:5, 1:5])
```

##	PM1:20080107	PM1:20080108	PM1:20080114
## Actinobacteria	0	3	2
## Bacteroidetes	486	921	1103
## Cyanobacteria	0	0	0
## Firmicutes	455	922	1637
## Proteobacteria	29	14	30
##	PM1:20071211	PM1:20080121	
## Actinobacteria	37	0	
## Bacteroidetes	607	818	
## Cyanobacteria	0	0	
## Firmicutes	772	1254	
## Proteobacteria	38	23	

The `aggregateByTaxonomy` and `aggTax` functions are flexible enough to put in either 1) a matrix with a vector of labels or 2) a MRexperiment object with a vector of labels or featureData column name. The function can also output either a matrix or MRexperiment object.

6 Visualization of features

Visualizing metagenomic datasets can be very revealing. To help with the analysis `metagenomeSeq` has several plotting functions to gain insight of the dataset's overall structure and particular individual features. We include an shiny based interactive first exploration of the data through the `display` function.

For an overall look at the dataset we provide a number of plots including heatmaps of feature counts: `plotMRheatmap`, basic feature correlation structures: `plotCorr`, PCA/MDS coordinates of samples or features: `plotOrd`, rarefaction effects: `plotRare` and contingency table style plots: `plotBubble`.

Other plotting functions look at particular features such as the abundance for a single feature: `plotOTU` and `plotFeature`, or of multiple features at once: `plotGenus`. Plotting multiple OTUs with similar annotations allows for additional control of false discoveries.

6.1 Interactive Display

Due to recent advances in the `interactiveDisplay` package, calling the `display` function on `MRExperiment` objects will bring up a browser to explore your data through several interactive visualizations. For more detailed interactive visualizations one might be interested in the `shiny-phyloseq` package.

```
# Calling display on the MRExperiment object will start a browser  
# session with interactive plots.  
  
# display(mouseData)
```

6.2 Structural overview

Many studies begin by comparing the abundance composition across sample or feature phenotypes. Often a first step of data analysis is a heatmap, correlation or co-occurrence plot or some other data exploratory method. The following functions have been implemented to provide a first step overview of the data:

1. `plotMRheatmap` - heatmap of abundance estimates (Fig. 4 right)
2. `plotCorr` - heatmap of pairwise correlations (Fig. 4 left)
3. `plotOrd` - PCA/CMDS components (Fig. 5 left)
4. `plotRare` - rarefaction effect (Fig. 5 right)
5. `plotBubble` - contingency table style plot (see help)

Each of the above can include phenotypic information in helping to explore the data.

Below we show an example of how to create a heatmap and hierarchical clustering of \log_2 transformed counts for the 200 OTUs with the largest overall variance. Red values indicate counts close to zero. Row color labels indicate OTU taxonomic class; column color labels indicate diet (green = high fat, yellow = low fat). Notice the samples cluster by diet in these cases and there are obvious clusters. We then plot a correlation matrix for the same features.

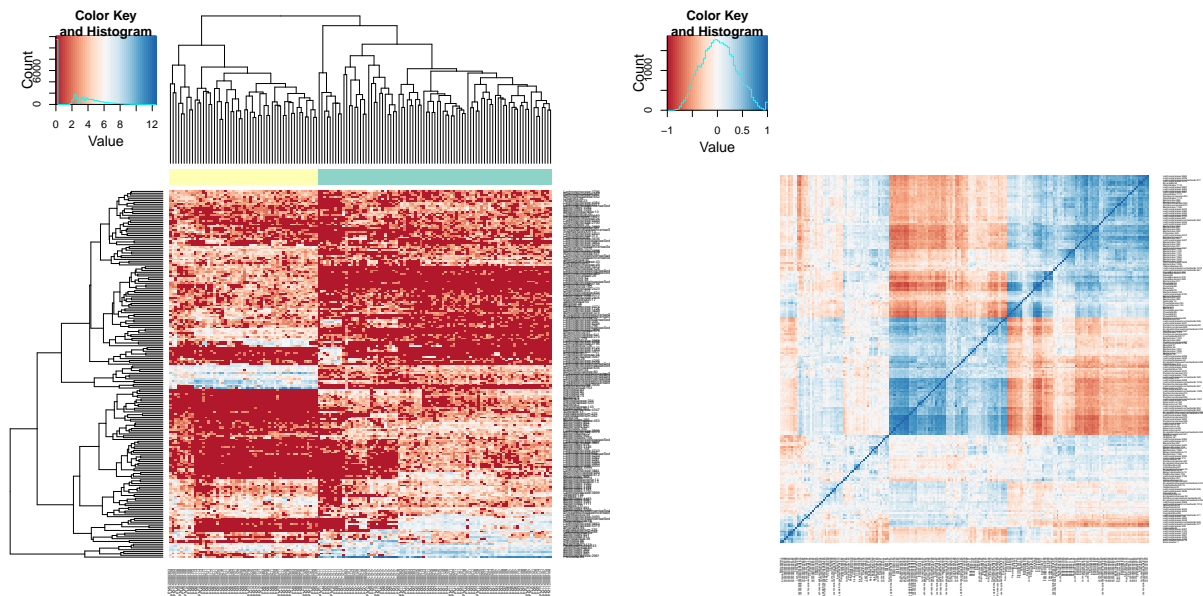


Figure 4: Left) Abundance heatmap (plotMRheatmap). Right) Correlation heatmap (plotCorr).

```

trials = pData(mouseData)$diet
heatmapColColors = brewer.pal(12, "Set3")[as.integer(factor(trials))]
heatmapCols = colorRampPalette(brewer.pal(9, "RdBu"))(50)

# plotMRheatmap
plotMRheatmap(obj = mouseData, n = 200, cexRow = 0.4, cexCol = 0.4,
               trace = "none", col = heatmapCols, ColSideColors = heatmapColColors)

# plotCorr
plotCorr(obj = mouseData, n = 200, cexRow = 0.25, cexCol = 0.25, trace = "none",
          dendrogram = "none", col = heatmapCols)

```

Below is an example of plotting CMDS plots of the data and the rarefaction effect at the OTU level. None of the data is removed (we recommend removing outliers typically).

```

cl = factor(pData(mouseData)$diet)

# plotOrd - can load vegan and set distfun = vegdist and use
# dist.method='bray'
plotOrd(mouseData, tran = TRUE, usePCA = FALSE, useDist = TRUE, bg = cl,
        pch = 21)

# plotRare
res = plotRare(mouseData, cl = cl, ret = TRUE, pch = 21, bg = cl)

# Linear fits for plotRare / legend
tmp = lapply(levels(cl), function(lv) lm(res[, "ident"] ~ res[, "libSize"] -
    1, subset = cl == lv))
for (i in 1:length(levels(cl))) {
  abline(tmp[[i]], col = i)
}

```

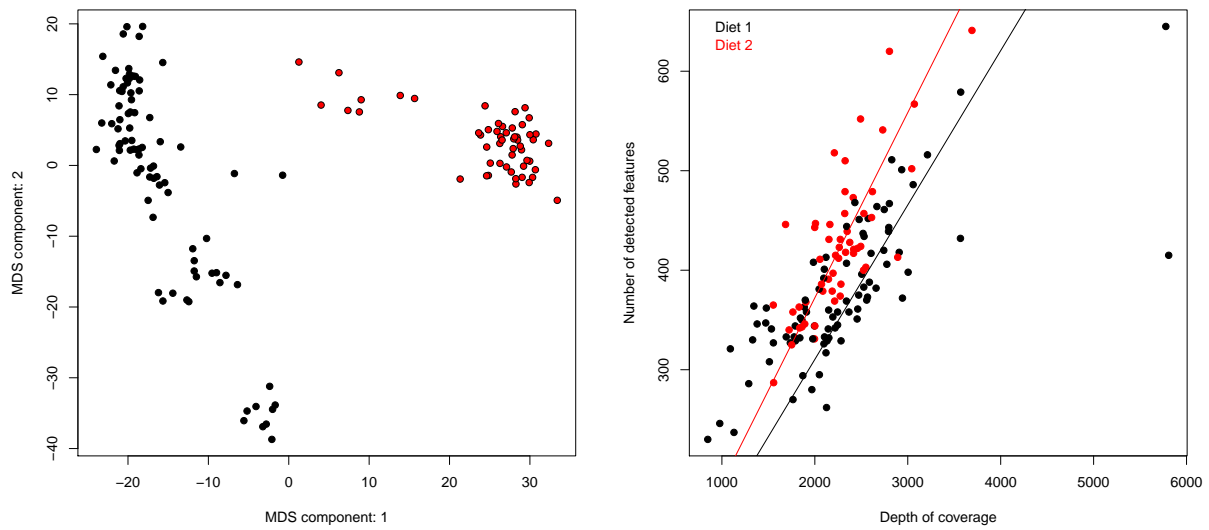



Figure 5: Left) CMDS of features (plotOrd). Right) Rarefaction effect (plotRare).

```
legend("topleft", c("Diet 1", "Diet 2"), text.col = c(1, 2), box.col = NA)
```

6.3 Feature specific

Reads clustered with high similarity represent functional or taxonomic units. However, it is possible that reads from the same organism get clustered into multiple OTUs. Following differential abundance analysis. It is important to confirm differential abundance. One way to limit false positives is ensure that the feature is actually abundant (enough positive samples). Another way is to plot the abundances of features similarly annotated.

1. plotOTU - abundances of a particular feature by group (Fig. 6 left)
2. plotGenus - abundances for several features similarly annotated by group (Fig. 6 right)
3. plotFeature - abundances of a particular feature by group (similar to plotOTU, Fig. 7)

Below we use plotOTU to plot the normalized log(cpt) of a specific OTU annotated as *Neisseria meningitidis*, in particular the 779th row of lungTrim's count matrix. Using plotGenus we plot the normalized log(cpt) of all OTUs annotated as *Neisseria meningitidis*.

It would appear that *Neisseria meningitidis* is differentially more abundant in nonsmokers.

```
head(MRtable(fit, coef = 2, taxa = 1:length(fData(lungTrim)$taxa)))
```

##	+samples in group 1	+samples in group 0	counts in group 1
## 63	6	24	11
## 779	7	23	22
## 358	1	24	1
## 499	2	21	2
## 25	26	15	1893
## 928	11	2	91

```
##      counts in group 0 smokingStatusSmoker      pvalues
## 63      1538      -4.031612 6.064532e-15
## 779      1512      -3.958899 1.617066e-14
## 358      390      -2.927686 1.869400e-12
## 499      326      -2.675306 8.705169e-10
## 25      162      2.575672 2.767250e-08
## 928      4      2.574172 1.382567e-37
##      adjPvalues
## 63 2.496161e-13
## 779 6.162817e-13
## 358 5.343369e-11
## 499 1.279660e-08
## 25 2.230904e-07
## 928 6.091832e-35

patients = sapply(strsplit(rownames(pData(lungTrim)), split = "_"),
  function(i) {
    i[3]
  })
pData(lungTrim)$patients = patients
classIndex = list(smoker = which(pData(lungTrim)$SmokingStatus ==
  "Smoker"))
classIndex$non smoker = which(pData(lungTrim)$SmokingStatus == "NonSmoker")
otu = 779

# plotOTU
plotOTU(lungTrim, otu = otu, classIndex, main = "Neisseria meningitidis")

# Now multiple OTUs annotated similarly
x = fData(lungTrim)$taxa[otu]
otulist = grep(x, fData(lungTrim)$taxa)

# plotGenus
plotGenus(lungTrim, otulist, classIndex, labs = FALSE, main = "Neisseria meningi

lablist <- c("S", "NS")
axis(1, at = seq(1, 6, by = 1), labels = rep(lablist, times = 3))

classIndex = list(Western = which(pData(mouseData)$diet == "Western"))
classIndex$BK = which(pData(mouseData)$diet == "BK")
otuIndex = 8770

# par(mfrow=c(1,2))
dates = pData(mouseData)$date
plotFeature(mouseData, norm = FALSE, log = FALSE, otuIndex, classIndex,
  col = dates, sortby = dates, ylab = "Raw reads")
```




Figure 6: Left) Abundance plot (plotOTU). Right) Multiple OTU abundances (plotGenus).

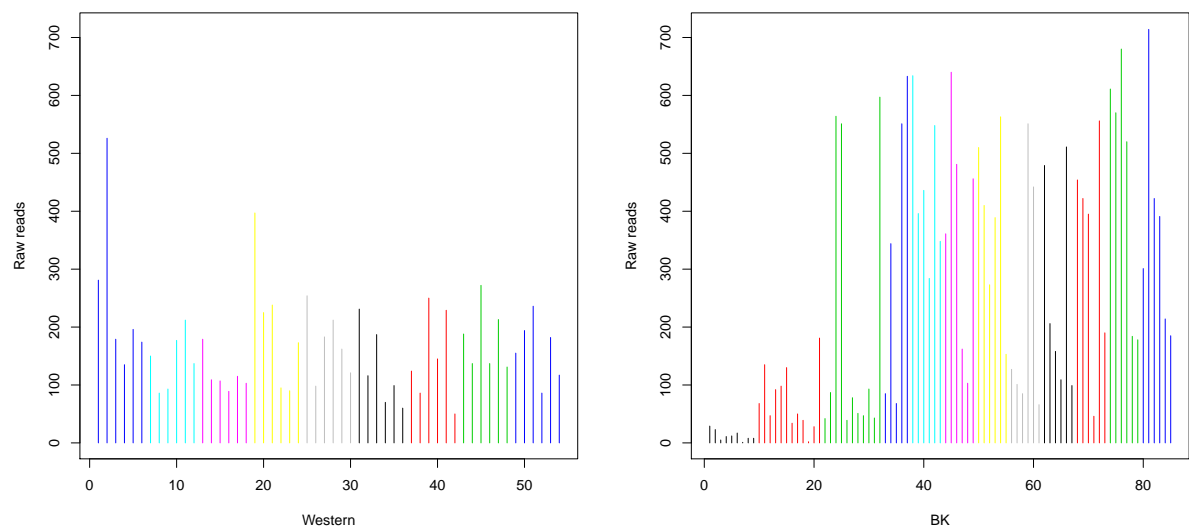


Figure 7: Plot of raw abundances

7 Summary

metagenomeSeq is specifically designed for sparse high-throughput sequencing experiments that addresses the analysis of differential abundance for marker gene survey data. The package, while designed for marker-gene survey datasets, may be appropriate for other sparse data sets for which the zero-inflated Gaussian mixture model may apply. If you make use of the statistical method please cite our paper. If you made use of the manual/software, please cite the manual/software!

7.1 Citing metagenomeSeq

```
citation("metagenomeSeq")

##
## Please cite the top for the original statistical method
## and normalization method implemented in metagenomeSeq and
## the bottom for the software/vignette guide. Time series
## analysis/function is described in the third citation.
##
## JN Paulson, OC Stine, HC Bravo, M Pop. Differential
## abundance analysis for microbial marker-gene surveys.
## Nat Meth Accepted
##
## JN Paulson, M Pop, HC Bravo. metagenomeSeq: Statistical
## analysis for sparse high-throughput sequencing.
## Bioconductor package: 1.8.3.
## http://cbcb.umd.edu/software/metagenomeSeq
##
## H Talukder, JN Paulson, HC Bravo. Finding regions of
## interest in high throughput genomics data using
## smoothing splines. Submitted
```

7.2 Session Info

```
sessionInfo()

## R version 3.1.2 (2014-10-31)
## Platform: x86_64-apple-darwin13.4.0 (64-bit)
##
## locale:
## [1] C/en_US.UTF-8/en_US.UTF-8/C/en_US.UTF-8/en_US.UTF-8
##
## attached base packages:
## [1] grid      parallel  stats      graphics  grDevices  utils
## [7] datasets  methods   base
##
## other attached packages:
## [1] gss_2.1-4          metagenomeSeq_1.8.3
## [3] RColorBrewer_1.1-2 interactiveDisplay_1.4.0
## [5] limma_3.22.1       Biobase_2.26.0
```

```
## [7] BiocGenerics_0.12.1      knitr_1.8
##
## loaded via a namespace (and not attached):
## [1] AnnotationDbi_1.28.1      Category_2.32.0
## [3] DBI_0.3.1                 GSEABase_1.28.0
## [5] GenomeInfoDb_1.2.4       IRanges_2.0.1
## [7] KernSmooth_2.23-13       MASS_7.3-35
## [9] Matrix_1.1-4             R.methodsS3_1.6.1
## [11] R6_2.0.1                 RBGL_1.42.0
## [13] RJSONIO_1.3-0            RSQLite_1.0.0
## [15] Rcpp_0.11.3              S4Vectors_0.4.0
## [17] XML_3.98-1.1             annotate_1.44.0
## [19] bitops_1.0-6             caTools_1.17.1
## [21] colorspace_1.2-4        digest_0.6.6
## [23] evaluate_0.5.5           formatR_1.0
## [25] gdata_2.13.3             genefilter_1.48.1
## [27] ggplot2_1.0.0            gplots_2.15.0
## [29] graph_1.44.1             gridSVG_1.4-2
## [31] gtable_0.1.2             gtools_3.4.1
## [33] highr_0.4               htmltools_0.2.6
## [35] httpuv_1.3.2             interactiveDisplayBase_1.4.0
## [37] lattice_0.20-29         matrixStats_0.12.2
## [39] mime_0.2                 munsell_0.4.2
## [41] plyr_1.8.1              proto_0.3-10
## [43] reshape2_1.4.1          scales_0.2.4
## [45] shiny_0.10.2.2          splines_3.1.2
## [47] stats4_3.1.2            stringr_0.6.2
## [49] survival_2.37-7         tools_3.1.2
## [51] xtable_1.7-4
```

8 Appendix

8.1 Appendix A: MRExperiment internals

The S4 class system in R allows for object oriented definitions. `metagenomeSeq` makes use of the `Biobase` package in Bioconductor and their virtual-class, `eSet`. Building off of `eSet`, the main S4 class in `metagenomeSeq` is termed `MRExperiment`. `MRExperiment` is a simple extension of `eSet`, adding a single slot, `expSummary`.

The experiment summary slot is a data frame that includes the depth of coverage and the normalization factors for each sample. Future datasets can be formatted as `MRExperiment` objects and analyzed with relative ease. A `MRExperiment` object is created by calling `newMRExperiment`, passing the counts, phenotype and feature data as parameters.

We do not include normalization factors or library size in the currently available slot specified for the sample specific phenotype data. All matrices are organized in the `assayData` slot. All phenotype data (disease status, age, etc.) is stored in `phenoData` and feature data (OTUs, taxonomic assignment to varying levels, etc.) in `featureData`. Additional slots are available for reproducibility and annotation.

8.2 Appendix B: Mathematical model

Defining the class comparison of interest as $k(j) = I\{j \in \text{group}A\}$. The zero-inflated model is defined for the continuity-corrected \log_2 of the count data $y_{ij} = \log_2(c_{ij} + 1)$ as a mixture of a point mass at zero $I_{\{0\}}(y_{ij})$ and a count distribution $f_{count}(y_{ij}; \mu_i, \sigma_i^2) \sim N(\mu_i, \sigma_i^2)$. Given mixture parameters π_j , we have that the density of the zero-inflated Gaussian distribution for feature i , in sample j with S_j total counts is:

$$f_{zig}(y_{ij}; \theta) = \pi_j(S_j) \cdot I_{\{0\}}(y_{ij}) + (1 - \pi_j(S_j)) \cdot f_{count}(y_{ij}; \theta) \quad (1)$$

Maximum-likelihood estimates are approximated using an EM algorithm, where we treat mixture membership $\Delta_{ij} = 1$ if y_{ij} is generated from the zero point mass as latent indicator variables [5]. We make use of an EM algorithm to account for the linear relationship between sparsity and depth of coverage. The user can specify within the `fitZig` function a non-default zero model that accounts for more than simply the depth of coverage (e.g. country, age, any metadata associated with sparsity, etc.). See Figure 8 for the graphical model.

More information will be included later. For now, please see the online methods in:

<http://www.nature.com/nmeth/journal/vaop/ncurrent/full/nmeth.2658.html>

8.3 Appendix C: Calculating the proper percentile

To be included: an overview of the two methods implemented for the data driven percentile calculation and more description below.

The choice of the appropriate quantile given is crucial for ensuring that the normalization approach does not introduce normalization-related artifacts in the data. At a high level, the count distribution of samples should all be roughly equivalent and independent of each other up to this quantile under the assumption that, at this range, counts are derived from a common distribution.

More information will be included later. For now, please see the online methods in:

<http://www.nature.com/nmeth/journal/vaop/ncurrent/full/nmeth.2658.html>

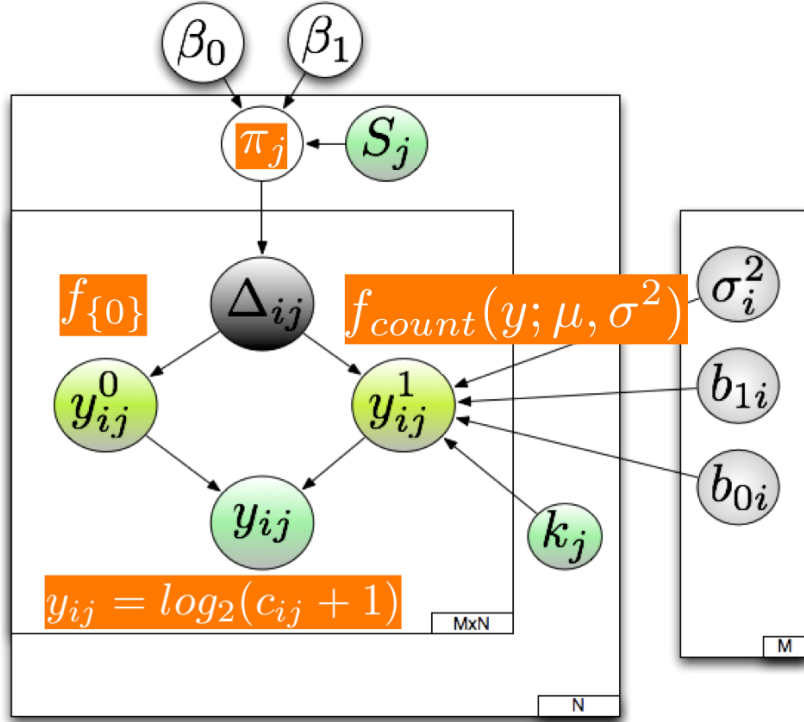


Figure 8: Graphical model. Green nodes represent observed variables: S_j is the total number of reads in sample j ; k_j the case-control status of sample j ; and y_{ij} the logged normalized counts for feature i in sample j . Yellow nodes represent counts obtained from each mixture component: counts come from either a spike-mass at zero, y_{ij}^0 , or the “count” distribution, y_{ij}^1 . Grey nodes b_{0i} , b_{1i} and σ_i^2 represent the estimated overall mean, fold-change and variance of the count distribution component for feature i . π_j , is the mixture proportion for sample j which depends on sequencing depth via a linear model defined by parameters β_0 and β_1 . The expected value of latent indicator variables Δ_{ij} give the posterior probability of a count being generated from a spike-mass at zero, i.e. y_{ij}^0 . We assume M features and N samples.

References

- [1] Emily S Charlson, Kyle Bittinger, Andrew R Haas, Ayannah S Fitzgerald, Ian Frank, Anjana Yadav, Frederic D Bushman, and Ronald G Collman. Topographical continuity of bacterial populations in the healthy human respiratory tract. *American Journal of Respiratory and Critical Care Medicine*, 184, 2011.
- [2] Peter J Turnbaugh, Vanessa K Ridaura, Jeremiah J Faith, Federico E Rey, Rob Knight, and Jeffrey I Gordon. The effect of diet on the human gut microbiome: a metagenomic analysis in humanized gnotobiotic mice. *Science translational medicine*, 1(6):6ra14, 2009.
- [3] Consortium HMP. A framework for human microbiome research. *Nature*, 486(7402), 2012.
- [4] Gordon K Smyth. *Limma: linear models for microarray data*. Number October. Springer, 2005.
- [5] A P Dempster, N M Laird, and D B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society Series B Methodological*, 39(1):1–38, 1977.