

# `goseq`: Gene Ontology testing for RNA-seq datasets

Matthew D. Young  
myoung@wehi.edu.au

Matthew J. Wakefield

Gordon K. Smyth

Alicia Oshlack

3 March 2011

## 1 Introduction

This document gives an introduction to the use of the `goseq` R Bioconductor package [Young et al., 2010]. This package provides methods for performing Gene Ontology analysis of RNA-seq data, taking length bias into account [Oshlack and Wakefield, 2009]. The methods and software used by `goseq` are equally applicable to other category based test of RNA-seq data, such as KEGG pathway analysis.

Once installed, the `goseq` package can be easily loaded into R using:

```
> library(goseq)
```

In order to perform a GO analysis of your RNA-seq data, `goseq` only requires a simple named vector, which contains two pieces of information.

1. **Measured genes:** all genes for which RNA-seq data was gathered for your experiment. Each element of your vector should be named by a unique gene identifier.
2. **Differentially expressed genes:** each element of your vector should be either a 1 or a 0, where 1 indicates that the gene is differentially expressed and 0 that it is not.

If the organism, gene identifier or category test is currently not natively supported by `goseq`, it will also be necessary to supply additional information regarding the genes length and/or the association between categories and genes.

A combination of bioconductor R packages such as `GenomicFeatures` and `Rsamtools` allow for the summarization of mapped reads into a table of counts, such as reads per gene. From there, several packages exist for performing differential expression analysis on summarized data

(eg. `edgeR` [Robinson and Smyth, 2007, 2008, Robinson et al., 2010]). `goseq` will work with any method for determining differential expression and as such differential expression analysis is outside the scope of this document, but in order to facilitate ease of use, we will make use of the `edgeR` package to calculate differentially expressed (DE) genes in all the case studies in this document.

## 2 Reading data

We assume that the user can use appropriate in-built R functions (such as `read.table` or `scan`) to obtain two vectors, one containing all genes assayed in the RNA-seq experiment, the other containing all genes which are DE. If we assume that the vector of genes being assayed is named `assayed.genes` and the vector of DE genes is named `de.genes` we can construct a named vector suitable for use with `goseq` using the following:

```
> gene.vector = as.integer(assayed.genes %in% de.genes)
> names(gene.vector) = assayed.genes
> head(gene.vector)
```

It may be that the user can already read in a vector in this format, in which case it can then be immediately used by `goseq`.

## 3 GO testing of RNA-seq data

To begin the analysis, `goseq` first needs to quantify the length bias present in the dataset under consideration. This is done by calculating a Probability Weighting Function or PWF which can be thought of as a function which gives the probability that a gene will be differentially expressed (DE), based on its length alone. The PWF is calculated by fitting a monotonic spline to the binary data series of differential expression (1=DE, 0=Not DE) as a function of gene length. The PWF is used to weight the chance of selecting each gene when forming a null distribution for GO category membership. The fact that the PWF is calculated directly from the dataset under consideration makes this approach robust, only correcting for the length bias present in the data. For example, if `goseq` is run on a microarray dataset, for which no length bias exists, the calculated PWF will be nearly flat and all genes will be weighted equally, resulting in no length bias correction.

In order to account for the length bias inherent to RNA-seq data when performing a GO analysis (or other category based tests), one cannot simply use the hypergeometric distribution as the null distribution for category membership, which is appropriate for data without DE length bias, such as microarray data. GO analysis of RNA-seq data requires the use of random sampling in order to generate a suitable null distribution for GO category membership and calculate each categories significance for over representation amongst DE genes.

However, this random sampling is computationally expensive. In most cases, the Wallenius distribution can be used to approximate the true null distribution, without any significant loss in

accuracy. The `goseq` package implements this approximation as its default option. The option to generate the null distribution using random sampling is also included as an option.

Having established a null distribution, each GO category is then tested for over and under representation amongst the set of differentially expressed genes and the null is used to calculate a p-value for under and over representation.

## 4 Natively supported Gene Identifiers and category tests

`goseq` needs to know the length of each gene, as well as what GO categories (or other categories of interest) each gene is associated with. `goseq` relies on the UCSC genome browser to provide the length information for each gene. However, because the process of fetching the length of every transcript is slow and bandwidth intensive, `goseq` relies on an offline copy of this information stored in the data package `geneLenDataBase`. To see which genome/gene identifier combinations are in the local database, simply run:

```
> supportedGenomes()
```

or

```
> supportedGeneIDs()
```

The rightmost column in the output of both these commands lists the identifiers (in the case of `supportedGenomes`) or genomes (in the case of `supportedGeneIDs`) for which length data exists in the local database. If your genome/ID combination is not in the local database, it will be downloaded from the UCSC genome browser. If your genome/ID combination is not in the UCSC database, you will have to manually specify the gene lengths.

In order to link GO categories to genes, `goseq` uses the organism packages from Bioconductor. These packages are named `org.<Genome>.<ID>.db`, where `<Genome>` is a short string identifying the genome and `<ID>` is a short string identifying the gene identifier. Currently, `goseq` will automatically retrieve the mapping between GO categories and genes from the relevant package (as long as it is installed) for commonly used genome/ID combinations. If GO mappings are not automatically available for your genome/ID combination, you will have to manually specify the relationship between genes and categories. Although the Genome/ID naming conventions used by the organism packages differ from the UCSC, `goseq` is able to convert between the two, so the user need only ever specify the UCSC genome/ID in most cases.

## 5 Non-native Gene Identifier or category test

If the organism, Gene Identifier or category test you wish to perform is not in the native `goseq` database, you will have to supply one or all of the following:

- **Length data:** the length of each gene in your gene identifier format.
- **Category mappings:** the mapping (usually many-to-many) between the categories you wish to test for over/under representation amongst DE genes and genes in your gene identifier format.

## 5.1 Length data format

The length data must be formatted as a numeric vector, of the same length as the main named vector specifying gene names/DE genes. Each entry should give the length of the corresponding gene in bp. If length data is unavailable for some genes, that entry should be set to NA.

## 5.2 Category mapping format

The mapping between category names and genes should be given as a data frame with two columns. One column should contain the gene IDs and the other the name of an associated category. As the mapping between categories and genes is usually many-to-many, this data frame will usually have multiple rows with the same gene name and category name.

Alternatively, mappings between genes and categories can be given as a list. The names of list entries should be gene IDs and the entries themselves should be a vector of category names to which the gene ID corresponds.

# 6 Case study: Prostate cancer data

## 6.1 Introduction

This section provides an analysis of data from an RNA-seq experiment to illustrate the use of `goseq` for GO analysis.

This experiment examined the effects of androgen stimulation on a human prostate cancer cell line, LNCaP. The data set includes more than 17 million short cDNA reads obtained for both the treated and untreated cell line and sequenced on Illumina's 1G genome analyzer.

For each sample we were provided with the raw 35 bp RNA-seq reads from the authors. For the untreated prostate cancer cells (LNCaP cell line) there were 4 lanes totaling 10 million, 35 bp reads. For the treated cells there were 3 lanes totaling 7 million, 35 bp reads. All replicates were technical replicates. Reads were mapped to NCBI version 36.3 of the human genome using bowtie. Any read with which mapped to multiple locations was discarded. Using the ENSEMBL 54 annotation from biomart, each mapped read was associated with an ENSEMBL gene. This was done by associating any read that overlapped with any part of the gene (not just the exons) with that gene. Reads that did not correspond to genes were discarded.

## 6.2 Source of the data

The data set used in this case study is taken from [Li et al., 2008] and was made available from the authors upon request.

## 6.3 Determining the DE genes using edgeR

To begin with, we load in the text data and convert it the appropriate edgeR DGEList object.

```
> library(edgeR)
> table.summary = read.table(system.file("extdata", "Li_sum.txt",
+   package = "goseq"), sep = "\t", header = TRUE, stringsAsFactors = FALSE)
> counts = table.summary[, -1]
> rownames(counts) = table.summary[, 1]
> grp = factor(rep(c("Control", "Treated"), times = c(4, 3)))
> summarized = DGEList(counts, lib.size = colSums(counts), group = grp)
```

Next, we use edgeR to estimate the biological dispersion and calculate differential expression using a negative binomial model.

```
> disp = estimateCommonDisp(summarized)
> disp$common.dispersion
```

```
[1] 0.05681933
```

```
> tested = exactTest(disp)
```

```
Comparison of groups: Treated - Control
```

```
> topTags(tested)
```

```
Comparison of groups: Treated-Control
```

|                 | logConc   | logFC     | PValue       | FDR          |
|-----------------|-----------|-----------|--------------|--------------|
| ENSG00000127954 | -31.02991 | 37.972297 | 6.800102e-79 | 3.366459e-74 |
| ENSG00000151503 | -12.96052 | 5.404687  | 9.143635e-66 | 2.263324e-61 |
| ENSG00000096060 | -11.77715 | 4.899691  | 4.866397e-60 | 8.030528e-56 |
| ENSG00000091879 | -15.35999 | 5.771018  | 5.469701e-55 | 6.769575e-51 |
| ENSG00000132437 | -14.14850 | -5.896416 | 3.487845e-52 | 3.453385e-48 |
| ENSG00000166451 | -12.61713 | 4.567569  | 4.410990e-52 | 3.626363e-48 |
| ENSG00000131016 | -14.80016 | 5.274233  | 5.127569e-52 | 3.626363e-48 |
| ENSG00000163492 | -17.28041 | 7.296034  | 1.231576e-44 | 7.621300e-41 |
| ENSG00000113594 | -12.24687 | 4.053165  | 6.002790e-44 | 3.301935e-40 |
| ENSG00000116285 | -13.01524 | 4.112220  | 4.005898e-43 | 1.983160e-39 |

Finally, we Format the DE genes into a vector suitable for use with `goseq`

```
> genes = as.integer(p.adjust(tested$table$p.value[tested$table$logFC !=
+ 0], method = "BH") < 0.05)
> names(genes) = row.names(tested$table[tested$table$logFC != 0, ])
> table(genes)
```

```
genes
  0    1
19344 3399
```

## 6.4 Determining Genome and Gene ID

In order to allow for automatic data retrieval, the user has to tell `goseq` what genome and gene ID format were used to summarize the data. In our case we used build 36.1 of the NCBI human genome, we check what code this corresponds to by running:

```
> head(supportedGenomes())
```

|   | db      | organism | date      | name                               |
|---|---------|----------|-----------|------------------------------------|
| 1 | hg19    | Human    | Feb. 2009 | Genome Reference Consortium GRCh37 |
| 2 | hg18    | Human    | Mar. 2006 | NCBI Build 36.1                    |
| 3 | hg17    | Human    | May 2004  | NCBI Build 35                      |
| 4 | hg16    | Human    | Jul. 2003 | NCBI Build 34                      |
| 5 | felCat4 | Cat      | Dec. 2008 | NHGRI catChrV17e                   |
| 6 | felCat3 | Cat      | Mar. 2006 | Broad Institute Release 3          |

```
1
2
3
4
5
6
ccdsGene,ensGene,exoniphy,geneSym
acembly,acescan,ccdsGene,ensGene,exoniphy,geneSymbol,geneid,genscan,knownGene,knownG
acembly,acescan,ccdsGene,ensGene,exoniphy,geneSymbol,geneid,genscan,knownGene,refGene,
acembly,ensGene,exoniphy,geneSymb
ensGene,geneSymbol,geneid,ge
```

Which lists the genome codes in the far left column, headed “db”. As we are using NCBI build 36.1, we see that we should use “hg18”. We also know that we used ENSEMBL Gene ID to summarize our read data, we check what code this corresponds to by running:

```
> head(supportedGeneIDs(), n = 12)
```

|   | db            | track          | subtrack | GeneID         |
|---|---------------|----------------|----------|----------------|
| 1 | knownGene     | UCSC Genes     | <NA>     | Entrez Gene ID |
| 2 | knownGeneOld3 | Old UCSC Genes | <NA>     |                |

```

3  wgEncodeGencodeManualV3  Gencode Genes      Gencode Manual HAVANA Pseudogene ID
4  wgEncodeGencodeAutoV3   Gencode Genes      Gencode Auto   HAVANA Pseudogene ID
5  wgEncodeGencodePolyaV3  Gencode Genes      Gencode PolyA  HAVANA Pseudogene ID
6  ccdsGene                 CCDS                <NA>
7  refGene                  RefSeq Genes       <NA>          Entrez Gene ID
8  xenoRefGene              Other RefSeq       <NA>
9  vegaGene                 Vega Genes         Vega Protein Genes HAVANA Pseudogene ID
10 vegaPseudoGene           Vega Genes         Vega Pseudogenes HAVANA Pseudogene ID
11 ensGene                  Ensembl Genes     <NA>          Ensembl gene ID
12 acembly                  AceView Genes     <NA>

```

```

1
2
3
4
5
6
7
8  anoCar1,aplCal1,braFlo1,caeJap1,caePb1,caePb2,caeRem2,caeRem3,calJac1,canFam1,canFam2
9
10
11
12

```

Note that the head command is only used here to save on space in this guide. Again, the gene ID codes are listed in the far left “db” column. We find that our gene ID code is “ensGene”. We will use these strings whenever we are asked for a genome or id.

## 6.5 GO analysis

### 6.5.1 Fitting the Probability Weighting Function (PWF)

We first need to obtain a weighting for each gene, depending on its length, given by the PWF. As you may have noticed when running supportedGenomes or supportedGeneIDs, length data is available in the local database for our gene ID, “ensGene” and our genome, “hg18”. We will let `goseq` automatically fetch this data from its databases.

```

> pwf = nullp(genes, "hg18", "ensGene")
> head(pwf)


```

```

          DEgenes bias.data      pwf
ENSG00000230758      0      NA      NA
ENSG00000182463      0    5455 0.2366278

```

|                 |   |      |           |
|-----------------|---|------|-----------|
| ENSG00000124208 | 0 | 2135 | 0.2020084 |
| ENSG00000230753 | 0 | NA   | NA        |
| ENSG00000224628 | 0 | NA   | NA        |
| ENSG00000125835 | 0 | 954  | 0.1509237 |



`nullp` plots the resulting fit, allowing verification of the goodness of fit before continuing the analysis. Further plotting of the pwf can be performed using the `plotPWF` function.

The output of `nullp` contains all the data used to create the PWF, as well as the PWF itself. It is a data frame with 3 columns, named "DEgenes", "bias.data" and "pwf" with the rownames set to the gene names. Each row corresponds to a gene with the DEgenes column specifying if the gene is DE (1 for DE, 0 for not DE), the bias.data column giving the numeric value of the DE bias being accounted for (usually the gene length or number of counts) and the pwf column giving the genes value on the probability weighting function.



### 6.5.2 Using the Wallenius approximation

To start with we will use the default method, to calculate the over and under expressed GO categories among DE genes. Again, we allow `goseq` to fetch data automatically, except this time the data being fetched is the relationship between ENSEMBL gene IDs and GO categories.

```
> GO.wall = goseq(pwf, "hg18", "ensGene")
> head(GO.wall)
```

|      | category   | over_represented_pvalue | under_represented_pvalue |
|------|------------|-------------------------|--------------------------|
| 2026 | GO:0005575 | 1.720902e-164           | 1                        |
| 2054 | GO:0005623 | 1.103968e-151           | 1                        |
| 8851 | GO:0044464 | 1.103968e-151           | 1                        |
| 1036 | GO:0003674 | 1.016040e-146           | 1                        |
| 3244 | GO:0008150 | 8.128373e-143           | 1                        |
| 3912 | GO:0009987 | 1.579216e-112           | 1                        |

The resulting object is ordered by GO category over representation amongst DE genes.

### 6.5.3 Using random sampling


It may sometimes be desirable to use random sampling to generate the null distribution for category membership. This is easily accomplished by using the `method` option to specify sampling and the `repcnt` option to specify the number of samples to generate:

```
> GO.samp = goseq(pwf, "hg18", "ensGene", method = "Sampling", repcnt = 1000)
> head(GO.samp)
```

|    | category   | over_represented_pvalue | under_represented_pvalue |
|----|------------|-------------------------|--------------------------|
| 2  | GO:0000003 | 0.000999001             | 1                        |
| 37 | GO:0000075 | 0.000999001             | 1                        |
| 40 | GO:0000079 | 0.000999001             | 1                        |
| 47 | GO:0000087 | 0.000999001             | 1                        |
| 75 | GO:0000139 | 0.000999001             | 1                        |
| 90 | GO:0000165 | 0.000999001             | 1                        |

You will notice that this takes far longer than the Wallenius approximation. Plotting the p-values against one another, we see that there is little difference between the two methods.

```
> plot(log10(GO.wall[, 2]), log10(GO.samp[match(GO.samp[, 1], GO.wall[,
+ 1]), 2]), xlab = "log10(Wallenius p-values)", ylab = "log10(Sampling p-values)",
+ xlim = c(-3, 0))
> abline(0, 1, col = 3, lty = 2)
```



#### 6.5.4 Ignoring length bias

`goseq` also allows for one to perform a GO analysis without correcting for RNA-seq length bias. In practice, this is only useful for assessing the effect of length bias on your results. You should NEVER use this option as your final analysis. If length bias is truly not present in your data, `goseq` will produce a nearly flat PWF and no length bias correction will be applied to your data and all methods will produce the same results.


However, if you still wish to ignore length bias in calculating GO category enrichment, this is again accomplished using the `method` option.

```
> GO.nobias = goseq(pwf, "hg18", "ensGene", method = "Hypergeometric")
> head(GO.nobias)
```

|      | category   | over_represented_pvalue | under_represented_pvalue |
|------|------------|-------------------------|--------------------------|
| 2026 | GO:0005575 | 6.220374e-208           | 1                        |
| 2054 | GO:0005623 | 6.182771e-193           | 1                        |
| 8851 | GO:0044464 | 6.182771e-193           | 1                        |
| 1036 | GO:0003674 | 4.572728e-190           | 1                        |
| 3244 | GO:0008150 | 2.283050e-181           | 1                        |
| 3912 | GO:0009987 | 4.721310e-145           | 1                        |

Ignoring length bias gives very different results from a length bias corrected analysis.

```
> plot(log10(GO.wall[, 2]), log10(GO.nobias[match(GO.nobias[, 1],
+   GO.wall[, 1]), 2]), xlab = "log10(Wallenius p-values)", ylab = "log10(Hypergeome
+   xlim = c(-3, 0), ylim = c(-3, 0))
> abline(0, 1, col = 3, lty = 2)
```



### 6.5.5 Limiting GO categories and other category based tests

By default, `goseq` tests all three major Gene Ontology branches; Cellular Components, Biological Processes and Molecular Functions. However, it is possible to limit testing to any combination of the major branches by using the `test.cats` argument to the `goseq` function. This is done by specifying a vector consisting of some combination of the strings “GO:CC”, “GO:BP” and “GO:MF”. For example, to test for only Molecular Function GO categories:

```
> GO.MF = goseq(pwf, "hg18", "ensGene", test.cats = c("GO:MF"))
> head(GO.MF)
```

|      | category   | over_represented_pvalue | under_represented_pvalue |   |
|------|------------|-------------------------|--------------------------|---|
| 123  | GO:0003674 | 1.016040e-146           |                          | 1 |
| 1077 | GO:0005488 | 5.565736e-100           |                          | 1 |
| 1087 | GO:0005515 | 4.446074e-70            |                          | 1 |
| 175  | GO:0003824 | 2.845162e-42            |                          | 1 |
| 19   | GO:0000166 | 8.224177e-20            |                          | 1 |
| 1935 | GO:0017076 | 2.399778e-18            |                          | 1 |

Native support for other category tests, such as KEGG pathway analysis are also made available via this argument. See the man `goseq` function man page for up to date information on what category tests are natively supported.

### 6.5.6 Making sense of the results

Having performed the GO analysis, you may now wish to interpret the results. If you wish to identify categories significantly enriched/unenriched below some p-value cutoff, it is necessary to first apply some kind of multiple hypothesis testing correction. For example, GO categories over enriched using a .05 FDR cutoff [Benjamini and Hochberg, 1995] are:

```
> enriched.GO = GO.wall$category[p.adjust(GO.wall$over_represented_pvalue,
+     method = "BH") < 0.05]
> head(enriched.GO)
```

```
[1] "GO:0005575" "GO:0005623" "GO:0044464" "GO:0003674" "GO:0008150" "GO:0009987"
```

Unless you are a machine, GO accession identifiers are probably not very meaningful to you. Information about each term can be obtained from the Gene Ontology website, <http://www.geneontology.org/>, or using the R package `GO.db`.

```
> library(GO.db)
> for (go in enriched.GO[1:10]) {
+   print(GOTERM[[go]])
+   cat("-----\n")
+ }
```

GOID: GO:0005575

Term: cellular\_component

Ontology: CC

Definition: The part of a cell or its extracellular environment in which a gene product is located. A gene product may be located in one or more parts of a cell and its location may be as specific as a particular macromolecular complex, that is, a stable, persistent association of macromolecules that function together.

Synonym: cellular component

Synonym: cellular component unknown

Synonym: GO:0008372

Secondary: GO:0008372

-----

GOID: GO:0005623

Term: cell

Ontology: CC

Definition: The basic structural and functional unit of all organisms. Includes the plasma membrane and any external encapsulating structures such as the cell wall and cell envelope.

-----

GOID: GO:0044464

Term: cell part

Ontology: CC

Definition: Any constituent part of a cell, the basic structural and functional unit of all organisms.

Synonym: protoplast

-----

GOID: GO:0003674

Term: molecular\_function

Ontology: MF

Definition: Elemental activities, such as catalysis or binding, describing the actions of a gene product at the molecular level. A given gene product may exhibit one or more molecular functions.

Synonym: molecular function

Synonym: molecular function unknown

Synonym: GO:0005554

Secondary: GO:0005554

-----

GOID: GO:0008150

Term: biological\_process

Ontology: BP

Definition: Any process specifically pertinent to the functioning of

integrated living units: cells, tissues, organs, and organisms. A process is a collection of molecular events with a defined beginning and end.

Synonym: biological process  
Synonym: biological process unknown  
Synonym: physiological process  
Synonym: GO:0000004  
Synonym: GO:0007582  
Secondary: GO:0000004  
Secondary: GO:0007582

-----  
GOID: GO:0009987  
Term: cellular process  
Ontology: BP  
Definition: Any process that is carried out at the cellular level, but not necessarily restricted to a single cell. For example, cell communication occurs among more than one cell, but occurs at the cellular level.  
Synonym: cell growth and/or maintenance  
Synonym: cell physiology  
Synonym: cellular physiological process  
Synonym: GO:0008151  
Synonym: GO:0050875  
Secondary: GO:0008151  
Secondary: GO:0050875

-----  
GOID: GO:0005488  
Term: binding  
Ontology: MF  
Definition: The selective, non-covalent, often stoichiometric, interaction of a molecule with one or more specific sites on another molecule.  
Synonym: ligand

-----  
GOID: GO:0005622  
Term: intracellular  
Ontology: CC  
Definition: The living contents of a cell; the matter contained within (but not including) the plasma membrane, usually taken to exclude large vacuoles and masses of secretory or ingested material. In eukaryotes it includes the nucleus and cytoplasm.  
Synonym: internal to cell  
Synonym: nucleocytoplasm

Synonym: protoplasm  
Synonym: protoplast

---

GOID: GO:0044424  
Term: intracellular part  
Ontology: CC  
Definition: Any constituent part of the living contents of a cell; the matter contained within (but not including) the plasma membrane, usually taken to exclude large vacuoles and masses of secretory or ingested material. In eukaryotes it includes the nucleus and cytoplasm.

---

GOID: GO:0005737  
Term: cytoplasm  
Ontology: CC  
Definition: All of the contents of a cell excluding the plasma membrane and nucleus, but including other subcellular structures.

---

### 6.5.7 Understanding `goseq` internals

The situation may arise where it is necessary for the user to perform some of the data processing steps usually performed automatically by `goseq` themselves. With this in mind, it will be useful to step through the preprocessing steps performed automatically by `goseq` to understand what is happening.

To start with, when `nullp` is called, `goseq` uses the genome and gene identifiers supplied to try and retrieve length information for all genes given to the `genes` argument. To do this, it retrieves the data from the database of gene lengths maintained in the package `geneLenDataBase`. This is performed by the `getlength` function in the following way:

```
> len = getlength(names(genes), "hg18", "ensGene")
> length(len)

[1] 22743

> length(genes)

[1] 22743

> head(len)

[1] NA 5455 2135 NA NA 954
```

After some data cleanup, the length data and the DE data is then passed to the `makespline` function to produce the PWF. The `nullp` returns a data frame which has 3 columns, the original DEgenes vector, the length bias data (in a column called `bias.data`) and the PWF itself (in a column named `pwf`). The names of the genes are also kept in this data frame as the names of the rows. If length data could not be obtained for a certain gene the corresponding entries in the "bias.data" and "pwf" columns are set to NA.

Next we call the `goseq` function to determine over/under representation of GO categories amongst DE genes. When we do this, `goseq` looks for the appropriate organism package and tries to obtain the mapping from genes to GO categories from it. This is done using the `getgo` function as follows:

```
> go = getgo(names(genes), "hg18", "ensGene")
> length(go)

[1] 22743

> length(genes)

[1] 22743

> head(go)

$<NA>
NULL

$ENSG00000182463
 [1] "GO:0006139" "GO:0006350" "GO:0006351" "GO:0006355" "GO:0006807" "GO:0007275"
 [7] "GO:0008150" "GO:0008152" "GO:0009058" "GO:0009059" "GO:0009889" "GO:0009987"
[13] "GO:0010467" "GO:0010468" "GO:0010556" "GO:0016070" "GO:0019219" "GO:0019222"
[19] "GO:0031323" "GO:0031326" "GO:0032501" "GO:0032502" "GO:0032774" "GO:0034641"
[25] "GO:0034645" "GO:0043170" "GO:0044237" "GO:0044238" "GO:0044249" "GO:0044260"
[31] "GO:0045449" "GO:0050789" "GO:0050794" "GO:0051171" "GO:0051252" "GO:0060255"
[37] "GO:0065007" "GO:0080090" "GO:0090304" "GO:0005575" "GO:0005622" "GO:0005623"
[43] "GO:0005634" "GO:0043226" "GO:0043227" "GO:0043229" "GO:0043231" "GO:0044424"
[49] "GO:0044464" "GO:0003674" "GO:0003676" "GO:0003677" "GO:0003700" "GO:0005488"
[55] "GO:0008270" "GO:0043167" "GO:0043169" "GO:0043565" "GO:0046872" "GO:0046914"

$ENSG00000124208
 [1] "GO:0000209" "GO:0006139" "GO:0006259" "GO:0006281" "GO:0006282" "GO:0006350"
 [7] "GO:0006351" "GO:0006355" "GO:0006464" "GO:0006807" "GO:0006950" "GO:0006974"
[13] "GO:0007154" "GO:0007165" "GO:0007243" "GO:0007249" "GO:0008150" "GO:0008152"
[19] "GO:0009058" "GO:0009059" "GO:0009889" "GO:0009966" "GO:0009967" "GO:0009987"
```



```

[25] "GO:0010467" "GO:0010468" "GO:0010556" "GO:0010627" "GO:0010646" "GO:0010647"
[31] "GO:0010740" "GO:0016070" "GO:0016567" "GO:0019219" "GO:0019222" "GO:0019538"
[37] "GO:0023014" "GO:0023033" "GO:0023034" "GO:0023046" "GO:0023051" "GO:0023052"
[43] "GO:0023056" "GO:0023060" "GO:0030154" "GO:0031323" "GO:0031326" "GO:0032446"
[49] "GO:0032502" "GO:0032774" "GO:0033554" "GO:0034641" "GO:0034645" "GO:0035466"
[55] "GO:0035468" "GO:0035556" "GO:0043122" "GO:0043123" "GO:0043170" "GO:0043388"
[61] "GO:0043412" "GO:0043687" "GO:0044093" "GO:0044237" "GO:0044238" "GO:0044249"
[67] "GO:0044260" "GO:0044267" "GO:0045449" "GO:0048518" "GO:0048522" "GO:0048583"
[73] "GO:0048869" "GO:0050789" "GO:0050794" "GO:0050896" "GO:0051052" "GO:0051090"
[79] "GO:0051091" "GO:0051092" "GO:0051098" "GO:0051099" "GO:0051101" "GO:0051171"
[85] "GO:0051246" "GO:0051252" "GO:0051716" "GO:0060255" "GO:0065007" "GO:0065009"
[91] "GO:0070647" "GO:0080090" "GO:0080134" "GO:0080135" "GO:0090046" "GO:0090047"
[97] "GO:0090304" "GO:0005575" "GO:0005622" "GO:0005623" "GO:0005634" "GO:0005737"
[103] "GO:0005829" "GO:0031371" "GO:0032991" "GO:0043226" "GO:0043227" "GO:0043229"
[109] "GO:0043231" "GO:0043234" "GO:0044424" "GO:0044444" "GO:0044464" "GO:0003674"
[115] "GO:0003824" "GO:0005488" "GO:0005515" "GO:0016563" "GO:0016874" "GO:0016879"
[121] "GO:0016881" "GO:0019787" "GO:0030528" "GO:0005783" "GO:0005789" "GO:0012505"
[127] "GO:0016020" "GO:0016021" "GO:0031090" "GO:0031224" "GO:0042175" "GO:0044422"
[133] "GO:0044425" "GO:0044432" "GO:0044446" "GO:0071212"

```

\$<NA>

NULL

\$<NA>

NULL

\$ENSG00000125835

```

[1] "GO:0000375" "GO:0000377" "GO:0000387" "GO:0000398" "GO:0006139" "GO:0006396"
[7] "GO:0006397" "GO:0006807" "GO:0008150" "GO:0008152" "GO:0008380" "GO:0009987"
[13] "GO:0010467" "GO:0016043" "GO:0016070" "GO:0016071" "GO:0022607" "GO:0022613"
[19] "GO:0022618" "GO:0034621" "GO:0034622" "GO:0034641" "GO:0043170" "GO:0043933"
[25] "GO:0044085" "GO:0044237" "GO:0044238" "GO:0044260" "GO:0065003" "GO:0090304"
[31] "GO:0005575" "GO:0005622" "GO:0005623" "GO:0005634" "GO:0005654" "GO:0005681"
[37] "GO:0005683" "GO:0005737" "GO:0005829" "GO:0030529" "GO:0030532" "GO:0031974"
[43] "GO:0031981" "GO:0032991" "GO:0043226" "GO:0043227" "GO:0043229" "GO:0043231"
[49] "GO:0043233" "GO:0044422" "GO:0044424" "GO:0044428" "GO:0044444" "GO:0044446"
[55] "GO:0044464" "GO:0070013" "GO:0003674" "GO:0003676" "GO:0003723" "GO:0005488"
[61] "GO:0005515"

```

Note that the first thing the `getgo` function does is to convert the UCSC genome/ID namings into the naming convention used by the organism packages. This is done using two hard coded conversion vectors that are included in the `goseq` package but usually hidden from the user.

```
> goseq:::ID_MAP
```

```
knownGene  refGene  ensGene  geneSymbol
"eg"       "eg"      "ENSEMBL" "SYMBOL"
```

```
> goseq:::ORG_PACKAGES
```

```
anoGam      Arabidopsis      bosTau      ce
"org.Ag.eg" "org.At.tair"   "org.Bt.eg" "org.Ce.eg"
canFam      dm              danRer      E. coli K12
"org.Cf.eg" "org.Dm.eg"     "org.Dr.eg" "org.EcK12.eg"
E. coli Sakai galGal         hg          mm
"org.EcSakai.eg" "org.Gg.eg"   "org.Hs.eg" "org.Mm.eg"
rheMac      Malaria        panTro      rn
"org.Mmu.eg" "org.Pf.plasmo" "org.Pt.eg" "org.Rn.eg"
sacCer      Pig            xenTro
"org.Sc.sgd" "org.Ss.eg"    "org.Xl.eg"
```

It is just as valid to run the length and GO category fetching as separate steps and then pass the result to the `nullp` and `goseq` functions using the `bias.data` and `gene2cat` arguments. Thus the following two blocks of code are equivalent:

```
> pwf = nullp(genes, "hg18", "ensGene")
> go = goseq(pwf, "hg18", "ensGene")
```

and

```
> gene_lengths = getlength(names(genes), "hg18", "ensGene")
> pwf = nullp(genes, bias.data = gene_lengths)
> go_map = getgo(names(genes), "hg18", "ensGene")
> go = goseq(pwf, "hg18", "ensGene", gene2cat = go_map)
```

## 6.6 KEGG pathway analysis

In order to illustrate performing a category test not present in the `goseq` database, we perform a KEGG pathway analysis. For human, the mapping from KEGG pathways to genes are stored in the package `org.Hs.eg.db`, in the object `org.Hs.egPATH`. In order to test for KEGG pathway over representation amongst DE genes, we need to extract this information and put it in a format that `goseq` understands. Unfortunately, the `org.Hs.eg.db` package does not contain direct mappings between ENSEMBL gene ID and KEGG pathway. Therefore, we have to construct this map by combining the ENSEMBL <-> Entrez and Entrez <-> KEGG mappings. This can be done using the following code:

```

> en2eg = as.list(org.Hs.egENSEMBL2EG)
> eg2kegg = as.list(org.Hs.egPATH)
> grepKEGG = function(id, mapkeys) {
+   unique(unlist(mapkeys[id], use.names = FALSE))
+ }
> kegg = lapply(en2eg, grepKEGG, eg2kegg)
> head(kegg)

```

Note that this step is quite time consuming. The code written here is not the most efficient way of producing this result, but the logic is much clearer than faster algorithms. The source code for `getgo` contains a more efficient routine.

We produce the PWF as before. Then, to perform a KEGG analysis, we simply make use of the `gene2cat` option in `goseq`:

```

> pwf = nullp(genes, "hg18", "ensGene")
> KEGG = goseq(pwf, gene2cat = kegg)
> head(KEGG)

```

Note that we do not have to tell the `goseq` function what organism and gene ID we are using as we are manually supplying the mapping between genes and categories.

KEGG analysis is shown as an illustration of how to supply your own mapping between gene ID and category, KEGG analysis is actually natively supported by `GOseq` and we could have performed it with the following code.

```

> pwf = nullp(genes, "hg18", "ensGene")
> KEGG = goseq(pwf, "hg18", "ensGene", test.cats = "KEGG")
> head(KEGG)

```

|     | category | over_represented_pvalue | under_represented_pvalue |
|-----|----------|-------------------------|--------------------------|
| 86  | 01100    | 3.825106e-15            | 1.0000000                |
| 88  | 03010    | 8.078326e-09            | 1.0000000                |
| 189 | 05200    | 4.886484e-08            | 1.0000000                |
| 163 | 04914    | 8.549522e-07            | 0.9999998                |
| 111 | 04115    | 1.271477e-06            | 0.9999997                |
| 77  | 00900    | 7.588112e-06            | 0.9999994                |

Noting that this time it was necessary to tell the `goseq` function that we are using HG18 and ENSEMBL gene ID, as the function needs this information to automatically construct the mapping from `geneid` to KEGG pathway.

## 6.7 Extracting mappings from organism packages

If you know that the information mapping gene ID to your categories of interest is contained in the organism packages, but `goseq` fails to fetch it automatically, you may want to extract it yourself and then pass it to the `goseq` function using the `gene2cat` argument. This is done in exactly the same way as extracting the KEGG to ENSEMBL mappings in the section “KEGG pathway analysis” above. This example is actually the worst case, where it is necessary to combine two mappings to get the desired list. If we had instead wanted the association between Entrez gene IDs and KEGG pathways, the following code would have been sufficient:

```
> kegg = as.list(org.Hs.egPATH)
> head(kegg)

$`1`
[1] NA

$`10`
[1] "00232" "00983" "01100"

$`100`
[1] "00230" "01100" "05340"

$`1000`
[1] "04514" "05412"

$`10000`
 [1] "04010" "04012" "04062" "04150" "04210" "04370" "04510" "04530" "04620" "04630"
[11] "04660" "04662" "04664" "04666" "04722" "04910" "04914" "04920" "05142" "05200"
[21] "05210" "05211" "05212" "05213" "05214" "05215" "05218" "05220" "05221" "05222"
[31] "05223"

$`100008586`
[1] NA
```

A note on fetching GO mappings from the organism. The data structure of GO is a directed acyclic graph. This means that in addition to each GO category being associated with a set of genes, it may also have children that are associated to other genes. It is important to use the `org.Hs.egGO2ALLEGS` and NOT the `org.Hs.egGO` object to create the mapping between GO categories and gene identifiers, as the latter does not include the links to genes arising from “child” GO categories. Thank you to Christopher Fjell for pointing this out.

## 6.8 Correcting for other biases

It is possible that in some circumstances you will wish to correct not just for length bias, but for the total number of counts. This can make sense because power to detect DE depends on the total number of counts a gene receives, which is the product of gene length and gene expression. So correcting for read count bias will compensate for all biases, known and unknown, in power to detect DE. On the other hand, it will also remove bias resulting from differences in expression level, which may not be desirable.

Correcting for count bias will produce a different PWF. Therefore, we need to tell `goseq` about the data on which the fraction DE depends when calculating the PWF using the `nullp` function. We then simply pass the result to `goseq` as usual.

So, in order to tell `goseq` to correct for read count bias instead of length bias, all you need to do is supply a numeric vector, containing the number of counts for each gene to `nullp`.

```
> countbias = rowSums(counts)[rowSums(counts) != 0]
> length(countbias)
```

```
[1] 22743
```

```
> length(genes)
```

```
[1] 22743
```

To use the count bias when doing GO analysis, simply pass this vector to `nullp` using the `bias.data` option. Note that we have to supply "hg18" and "ensGene" to `goseq` as it is not used by `nullp` and hence not in the `pwf.counts` object.

```
> pwf.counts = nullp(genes, bias.data = countbias)
> GO.counts = goseq(pwf.counts, "hg18", "ensGene")
> head(GO.counts)
```

|      | category   | over_represented_pvalue | under_represented_pvalue |
|------|------------|-------------------------|--------------------------|
| 2026 | GO:0005575 | 2.087788e-74            | 1                        |
| 3244 | GO:0008150 | 1.278708e-63            | 1                        |
| 2054 | GO:0005623 | 1.429128e-63            | 1                        |
| 8851 | GO:0044464 | 1.429128e-63            | 1                        |
| 1036 | GO:0003674 | 4.128580e-63            | 1                        |
| 3912 | GO:0009987 | 9.753224e-46            | 1                        |

Note that if you want to correct for length bias, but your organism/gene identifier is not natively supported, then you need to follow the same procedure as above, only the numeric vector supplied will contain each gene's length instead of its number of reads.

## 7 Setup

This vignette was built on:

```
> sessionInfo()
```

```
R version 2.12.2 (2011-02-25)
```

```
Platform: x86_64-unknown-linux-gnu (64-bit)
```

```
locale:
```

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

```
attached base packages:
```

```
[1] stats      graphics  grDevices  utils      datasets  methods   base
```

```
other attached packages:
```

```
[1] GO.db_2.4.5           org.Hs.eg.db_2.4.6     RSQLite_0.9-4
[4] DBI_0.2-5             AnnotationDbi_1.12.1   Biobase_2.10.0
[7] edgeR_2.0.5          goseq_1.2.1           geneLenDataBase_0.99.5
[10] BiasedUrn_1.03
```

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

```
[1] BSgenome_1.18.3      Biostrings_2.18.4     GenomicFeatures_1.2.5
[4] GenomicRanges_1.2.3 IRanges_1.8.9         Matrix_0.999375-49
[7] RCurl_1.5-0          XML_3.2-0             biomaRt_2.6.0
[10] grid_2.12.2          lattice_0.19-23      limma_3.6.9
[13] mgcv_1.7-5           nlme_3.1-98          rtracklayer_1.10.6
[16] tools_2.12.2
```

## 8 Acknowledgments

Christopher Fjell for a series of bug fixes and pointing out the difference between the egGO and egGO2ALLEGS objects in the organism packages.

## References

- Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57:289–300, 1995.
- Hairi Li, Michael T Lovci, Young-Soo Kwon, Michael G Rosenfeld, Xiang-Dong Fu, and Gene W Yeo. Determination of tag density required for digital transcriptome analysis: application to an androgen-sensitive prostate cancer model. *Proc Natl Acad Sci USA*, 105(51):20179–84, Dec 2008. doi: 10.1073/pnas.0807121105.
- Alicia Oshlack and Matthew J Wakefield. Transcript length bias in RNA-seq data confounds systems biology. *Biol Direct*, 4:14, Jan 2009. doi: 10.1186/1745-6150-4-14.
- M. D. Robinson and G. K. Smyth. Moderated statistical tests for assessing differences in tag abundance. *Bioinformatics*, 23(21):2881–2887, 2007.
- M. D. Robinson and G. K. Smyth. Small-sample estimation of negative binomial dispersion, with applications to sage data. *Biostatistics*, 9(2):321–332, 2008.
- M. D. Robinson, D. J. McCarthy, and G. K. Smyth. edgeR: a bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, 26(1):139–40, Jan 2010. doi: 10.1093/bioinformatics/btp616.
- M. D. Young, M. J. Wakefield, G. K. Smyth, and A. Oshlack. Gene ontology analysis for RNA-seq: accounting for selection bias. *Genome Biology*, 11:R14, Feb 2010.