

How To Use GOstats and Category to do Hypergeometric testing with unsupported model organisms

M. Carlson

May 11, 2023

1 Introduction

This vignette is meant as an extension of what already exists in the GOstatsHyperG.pdf vignette. It is intended to explain how a user can run hypergeometric testing on GO terms or KEGG terms when the organism in question is not supported by an annotation package. The 1st thing for a user to determine then is whether or not their organism is supported by an organism package through AnnotationForge. In order to do that, they need only to call the `available.dbschemas()` method from AnnotationForge.

```
> library("AnnotationForge")
> available.dbschemas()

[1] "ARABIDOPSISCHIP_DB" "BOVINECHIP_DB"      "BOVINE_DB"
[4] "CANINECHIP_DB"      "CANINE_DB"          "CHICKENCHIP_DB"
[7] "CHICKEN_DB"         "ECOLICHIP_DB"       "ECOLI_DB"
[10] "FLYCHIP_DB"         "FLY_DB"             "GO_DB"
[13] "HUMANCHIP_DB"       "HUMANCROSSCHIP_DB"  "HUMAN_DB"
[16] "INPARANOIDDROME_DB" "INPARANOIDHOMSA_DB" "INPARANOIDMUSMU_DB"
[19] "INPARANOIDRATNO_DB" "INPARANOIDSACCE_DB" "KEGG_DB"
[22] "MALARIA_DB"         "MOUSECHIP_DB"       "MOUSE_DB"
[25] "PFAM_DB"            "PIGCHIP_DB"         "PIG_DB"
[28] "RATCHIP_DB"         "RAT_DB"             "WORMCHIP_DB"
[31] "WORM_DB"            "YEASTCHIP_DB"       "YEAST_DB"
[34] "ZEBRAFISHCHIP_DB"   "ZEBRAFISH_DB"
```

If the organism that you are using is listed here, then your organism is supported. If not, then you will need to find a source or GO (or KEGG) to gene mappings. One source for GO to gene mappings is the blast2GO project. But you might also find such mappings at Ensembl or NCBI. If your organism is not a typical model organism, then the GO terms you will find are probably going to be predictions based on sequence similarity measures instead of direct measurements. This is something that you might want to bear in mind when you draw conclusions later.

In preparation for our subsequent demonstrations, let's get some data to work with by borrowing from an organism package. We will assume that you will use something like `read.table` to load your own annotation packages into a data.frame object. The starting object needs to be a data.frame with the GO Id's in the 1st col, the evidence codes in the 2nd column and the gene Id's in the 3rd.

```
> library("org.Hs.eg.db")
> frame = toTable(org.Hs.egGO)
> goframeData = data.frame(frame$go_id, frame$Evidence, frame$gene_id)
> head(goframeData)
```

```
   frame.go_id frame.Evidence frame.gene_id
1   GO:0008150             ND             1
2   GO:0001553             IEA             2
3   GO:0001869             IDA             2
4   GO:0002438             IEA             2
5   GO:0006953             IEA             2
6   GO:0007584             IEA             2
```

1.1 Preparing GO to gene mappings

When using GO mappings, it is important to consider the data structure of the GO ontologies. The terms are organized into a directed acyclic graph. The structure of the graph creates implications about the mappings of less specific terms to genes that are mapped to more specific terms. The Category and GOstats packages normally deal with this kind of complexity by using a special GO2All mapping in the annotation packages. You won't have one of those, so instead you will have to make one. AnnotationDbi provides some simple tools to represent your GO term to gene mappings so that this process will be easy. First you need to put your data into a GOFrame object. Then the simple act of casting this object to a GOAllFrame object will tap into the GO.db package and populate this object with the implicated GO2All mappings for you.

```
> goFrame=GOFrame(goframeData,organism="Homo sapiens")
> goAllFrame=GOAllFrame(goFrame)
```

In an effort to standardize the way that we pass this kind of custom information around, we have chosen to support geneSetCollection objects from GSEABase package. You can generate one of these objects in the following way:

```
> library("GSEABase")
> gsc <- GeneSetCollection(goAllFrame, setType = GOCollection())
```

1.2 Setting up the parameter object

Now we can make a parameter object for GOstats by using a special constructor function. For the sake of demonstration, I will just use all the EGs as the universe and grab some arbitrarily to be the geneIds tested. For your case, you need to make sure that the gene IDs you use are unique and that they are the same type for both the universe, the geneIds and the IDs that are part of your geneSetCollection.

```
> library("GOstats")
> universe = Lkeys(org.Hs.egGO)
> genes = universe[1:500]
> params <- GSEAGOHyperGParams(name="My Custom GSEA based annot Params",
```

```

+                                     geneSetCollection=gsc,
+                                     geneIds = genes,
+                                     universeGeneIds = universe,
+                                     ontology = "MF",
+                                     pvalueCutoff = 0.05,
+                                     conditional = FALSE,
+                                     testDirection = "over")

```

And finally we can call `hyperGTest` in the same way we always have before.

```

> Over <- hyperGTest(params)
> head(summary(Over))

```

	GOMFID	Pvalue	OddsRatio	ExpCount	Count	Size	
1	GO:0042626	1.699575e-32	25.40483	2.2635963	35	99	
2	GO:0019829	1.222928e-30	47.35966	1.2118243	27	53	
3	GO:0003824	2.103362e-26	2.88750	126.4870162	230	5532	
4	GO:0015662	9.910448e-26	252.34243	0.4572922	17	20	
5	GO:0015399	1.433430e-23	11.99596	3.8869835	35	170	
6	GO:0140358	2.766880e-23	108.12265	0.5487506	17	24	
							Term
1							ATPase-coupled transmembrane transporter activity
2							ATPase-coupled monoatomic cation transmembrane transporter activity
3							catalytic activity
4							P-type ion transporter activity
5							primary active transmembrane transporter activity
6							P-type transmembrane transporter activity

1.3 Preparing KEGG to gene mappings

This is much the same as what you did with the GO mappings except for two important simplifications. First of all you will no longer need to track evidence codes, so the object you start with only needs to hold KEGG Ids and gene IDs. Secondly, since KEGG is not a directed graph, there is no need for a KEGG to All mapping. Once again I will borrow some data to use as an example. Notice that we have to put the KEGG Ids in the left hand column of our initial two column data.frame.

```

> frame = toTable(org.Hs.egPATH)
> keggframeData = data.frame(frame$path_id, frame$gene_id)
> head(keggframeData)

```

	frame.path_id	frame.gene_id
1	04610	2
2	00232	9
3	00983	9
4	01100	9
5	00232	10
6	00983	10

```
> keggFrame=KEGGFrame(keggframeData,organism="Homo sapiens")
```

The rest of the process should be very similar.

```
> gsc <- GeneSetCollection(keggFrame, setType = KEGGCollection())
> universe = Lkeys(org.Hs.egGO)
> genes = universe[1:500]
> kparams <- GSEAKEGGHyperGParams(name="My Custom GSEA based annot Params",
+                                 geneSetCollection=gsc,
+                                 geneIds = genes,
+                                 universeGeneIds = universe,
+                                 pvalueCutoff = 0.05,
+                                 testDirection = "over")
> kOver <- hyperGTest(params)
> head(summary(kOver))
```

	GOMFID	Pvalue	OddsRatio	ExpCount	Count	Size	
1	GO:0042626	1.699575e-32	25.40483	2.2635963	35	99	
2	GO:0019829	1.222928e-30	47.35966	1.2118243	27	53	
3	GO:0003824	2.103362e-26	2.88750	126.4870162	230	5532	
4	GO:0015662	9.910448e-26	252.34243	0.4572922	17	20	
5	GO:0015399	1.433430e-23	11.99596	3.8869835	35	170	
6	GO:0140358	2.766880e-23	108.12265	0.5487506	17	24	
							Term
1							ATPase-coupled transmembrane transporter activity
2							ATPase-coupled monoatomic cation transmembrane transporter activity
3							catalytic activity
4							P-type ion transporter activity
5							primary active transmembrane transporter activity
6							P-type transmembrane transporter activity

```
> toLatex(sessionInfo())
```

- R version 4.3.0 RC (2023-04-13 r84257), x86_64-apple-darwin20
- Locale: C/en_US.UTF-8/en_US.UTF-8/C/en_US.UTF-8/en_US.UTF-8
- Time zone: America/New_York
- TZcode source: internal
- Running under: macOS Monterey 12.6.4
- Matrix products: default
- BLAS:
/Library/Frameworks/R.framework/Versions/4.3-x86_64/Resources/lib/libRblas.0.

- LAPACK:
`/Library/Frameworks/R.framework/Versions/4.3-x86_64/Resources/lib/libRlapack.`
`; LAPACK version3.11.0`
- Base packages: base, datasets, grDevices, graphics, grid, methods, stats, stats4, utils
- Other packages: ALL 1.42.0, AnnotationDbi 1.62.1, AnnotationForge 1.42.0, Biobase 2.60.0, BiocGenerics 0.46.0, BiocStyle 2.28.0, Category 2.66.0, GO.db 3.17.0, GOSTats 2.66.0, GSEABase 1.62.0, IRanges 2.34.0, Matrix 1.5-4, RBGL 1.76.0, RColorBrewer 1.1-3, Rgraphviz 2.44.0, S4Vectors 0.38.1, XML 3.99-0.14, annotate 1.78.0, genefilter 1.82.1, geneplotter 1.78.0, graph 1.78.0, hgu95av2.db 3.13.0, lattice 0.21-8, multtest 2.56.0, org.Hs.eg.db 3.17.0, xtable 1.8-4
- Loaded via a namespace (and not attached): BiocManager 1.30.20, Biostrings 2.68.0, DBI 1.1.3, GenomeInfoDb 1.36.0, GenomeInfoDbData 1.2.10, KEGGREST 1.40.0, MASS 7.3-60, MatrixGenerics 1.12.0, R6 2.5.1, RCurl 1.98-1.12, RSQLite 2.3.1, Rcpp 1.0.10, XVector 0.40.0, bit 4.0.5, bit64 4.0.5, bitops 1.0-7, blob 1.2.4, bookdown 0.34, bslib 0.4.2, cachem 1.0.8, cli 3.6.1, codetools 0.2-19, compiler 4.3.0, crayon 1.5.2, curl 5.0.0, digest 0.6.31, evaluate 0.21, fastmap 1.1.1, highr 0.10, htmltools 0.5.5, httr 1.4.6, jquerylib 0.1.4, jsonlite 1.8.4, knitr 1.42, magick 2.7.4, magrittr 2.0.3, matrixStats 0.63.0, memoise 2.0.1, pkgconfig 2.0.3, png 0.1-8, rlang 1.1.1, rmarkdown 2.21, sass 0.4.6, splines 4.3.0, survival 3.5-5, tools 4.3.0, vctrs 0.6.2, xfun 0.39, yaml 2.3.7, zlibbioc 1.46.0

>