

AnnotationDbi: Introduction To Bioconductor Annotation Packages

Marc Carlson

April 5, 2013

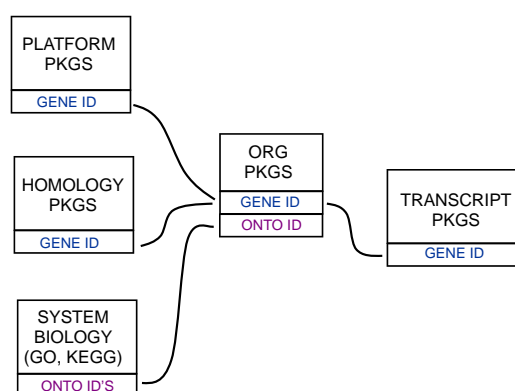


Figure 1: Annotation Packages: the big picture

Bioconductor provides extensive annotation resources. These can be *gene centric*, or *genome centric*. Annotations can be provided in packages curated by *Bioconductor*, or obtained from web-based resources. This vignette is primarily concerned with describing the annotation resources that are available as packages. More advanced users who wish to learn about how to make new annotation packages should see the vignette titled "Creating select Interfaces for custom Annotation resources" from the *AnnotationForge* package.

Gene centric *AnnotationDbi* packages include:

- Organism level: e.g. *org.Mm.eg.db*.

- Platform level: e.g. *hgu133plus2.db*, *hgu133plus2.probes*, *hgu133plus2.cdf*.
- Homology level: e.g. *hom.Dm.inp.db*.
- System-biology level: *GO.db*

Genome centric *GenomicFeatures* packages include

- Transcriptome level: e.g. *TxDb.Hsapiens.UCSC.hg19.knownGene*
- Generic genome features: Can generate via *GenomicFeatures*

One web-based resource accesses [biomart](#), via the *biomaRt* package:

- Query web-based ‘biomart’ resource for genes, sequence, SNPs, and etc.

The most popular annotation packages have been modified so that they can make use of a new set of methods to more easily access their contents. These four methods are named: **cols**, **keytypes**, **keys** and **select**. And they are described in this vignette. They can currently be used with all chip, organism, and *TranscriptDb* packages along with the popular *GO.db* package.

For the older less popular packages, there are still convenient ways to retrieve the data. The *How to use bimap from the ".db" annotation packages* vignette in the *AnnotationDbi* package is a key reference for learn about how to use bimap objects.

Finally, all of the ‘.db’ (and most other *Bioconductor* annotation packages) are updated every 6 months corresponding to each release of *Bioconductor*. Exceptions are made for packages where the actual resources that the packages are based on have not themselves been updated.

0.1 AnnotationDb objects and the select method

As previously mentioned, a new set of methods have been added that allow a simpler way of extracting identifier based annotations. All the annotation packages that support these new methods expose an object named exactly the same as the package itself. These objects are collectively called *AnnotationDb* objects for the class that they all inherit from. The more specific classes (the ones that you will actually see in the wild) have names like *OrgDb*, *ChipDb* or *TranscriptDb* objects. These names correspond to the kind of package (and underlying schema) being represented. The methods that can be applied to all of these objects are **cols**, **keys**, **keytypes** and **select**.

0.2 ChipDb objects and the select method

An extremely common kind of Annotation package is the so called platform based or chip based package type. This package is intended to make the manufacturer labels for a series of probes or probesets to a wide range of gene-based features. A package of this kind will load an *ChipDb* object. Below is a set of examples to show how you might use the standard 4 methods to interact with an object of this type.

First we need to load the package:

```
R> library(hgu95av2.db)
```

If we list the contents of this package, we can see that one of the many things loaded is an object named after the package "hgu95av2.db":

```
R> ls("package:hgu95av2.db")

[1] "hgu95av2"                "hgu95av2.db"
[3] "hgu95av2ACCNUM"          "hgu95av2ALIAS2PROBE"
[5] "hgu95av2CHR"            "hgu95av2CHRLNGTHS"
[7] "hgu95av2CHRLOC"         "hgu95av2CHRLOCEND"
[9] "hgu95av2ENSEMBL"        "hgu95av2ENSEMBL2PROBE"
[11] "hgu95av2ENTREZID"       "hgu95av2ENZYME"
[13] "hgu95av2ENZYME2PROBE"   "hgu95av2GENENAME"
[15] "hgu95av2GO"             "hgu95av2GO2ALLPROBES"
[17] "hgu95av2GO2PROBE"       "hgu95av2MAP"
[19] "hgu95av2MAPCOUNTS"     "hgu95av2OMIM"
[21] "hgu95av2ORGANISM"       "hgu95av2ORGPKG"
[23] "hgu95av2PATH"           "hgu95av2PATH2PROBE"
[25] "hgu95av2PFAM"           "hgu95av2PMID"
[27] "hgu95av2PMID2PROBE"     "hgu95av2PROSITE"
[29] "hgu95av2REFSEQ"         "hgu95av2SYMBOL"
[31] "hgu95av2UNIGENE"        "hgu95av2UNIPROT"
[33] "hgu95av2_dbInfo"        "hgu95av2_dbconn"
[35] "hgu95av2_dbfile"        "hgu95av2_dbschema"
```

We can look at this object to learn more about it:

```
R> hgu95av2.db
```

ChipDb object:

```
| DBSCHEMAVERSION: 2.1
```

```

| Db type: ChipDb
| Supporting package: AnnotationDbi
| DBSCHEMA: HUMANCHIP_DB
| ORGANISM: Homo sapiens
| SPECIES: Human
| MANUFACTURER: Affymetrix
| CHIPNAME: Human Genome U95 Set
| MANUFACTURERURL: http://www.affymetrix.com/support/technical/byproduct.affx?product
| EGSOURCEDATE: 2013-Mar5
| EGSOURCENAME: Entrez Gene
| EGSOURCEURL: ftp://ftp.ncbi.nlm.nih.gov/gene/DATA
| CENTRALID: ENTREZID
| TAXID: 9606
| GOSOURCENAME: Gene Ontology
| GOSOURCEURL: ftp://ftp.geneontology.org/pub/go/godatabase/archive/latest-lite/
| GOSOURCEDATE: 20130302
| GOEGSOURCEDATE: 2013-Mar5
| GOEGSOURCENAME: Entrez Gene
| GOEGSOURCEURL: ftp://ftp.ncbi.nlm.nih.gov/gene/DATA
| KEGGSOURCENAME: KEGG GENOME
| KEGGSOURCEURL: ftp://ftp.genome.jp/pub/kegg/genomes
| KEGGSOURCEDATE: 2011-Mar15
| GPSOURCENAME: UCSC Genome Bioinformatics (Homo sapiens)
| GPSOURCEURL: ftp://hgdownload.cse.ucsc.edu/goldenPath/hg19
| GPSOURCEDATE: 2010-Mar22
| ENSOURCEDATE: 2013-Jan16
| ENSOURCENAME: Ensembl
| ENSOURCEURL: ftp://ftp.ensembl.org/pub/current_fasta
| UPSOURCENAME: Uniprot
| UPSOURCEURL: http://www.UniProt.org/
| UPSOURCEDATE: Thu Mar 7 17:01:43 2013

```

If we want to know what kinds of data are retrieveable via `select`, then we should use the `cols` method like this:

```
R> cols(hgu95av2.db)
```

[1]	"PROBEID"	"ENTREZID"	"PFAM"	"IPI"
[5]	"PROSITE"	"ACCNUM"	"ALIAS"	"CHR"
[9]	"CHRLOC"	"CHRLOCEND"	"ENZYME"	"MAP"

[13]	"PATH"	"PMID"	"REFSEQ"	"SYMBOL"
[17]	"UNIGENE"	"ENSEMBL"	"ENSEMBLPROT"	"ENSEMBLTRANS"
[21]	"GENENAME"	"UNIPROT"	"GO"	"EVIDENCE"
[25]	"ONTOLOGY"	"GOALL"	"EVIDENCEALL"	"ONTOLOGYALL"
[29]	"OMIM"	"UCSCKG"		

If we are further curious to know more about those values for cols, we can consult the help pages. Asking about any of these values will pull up a manual page describing the different fields and what they mean.

```
R> help("SYMBOL")
```

If we are curious about what kinds of fields we could potentially use as keys to query the database, we can use the `keytypes` method. In a perfect world, this method will return values very similar to what was returned by `cols`, but in reality, some kinds of values make poor keys and so this list is often shorter.

```
R> keytypes(hgu95av2.db)
```

[1]	"ENTREZID"	"PFAM"	"IPI"	"PROSITE"
[5]	"ACCNUM"	"ALIAS"	"CHR"	"CHRLOC"
[9]	"CHRLOCEND"	"ENZYME"	"MAP"	"PATH"
[13]	"PMID"	"REFSEQ"	"SYMBOL"	"UNIGENE"
[17]	"ENSEMBL"	"ENSEMBLPROT"	"ENSEMBLTRANS"	"GENENAME"
[21]	"UNIPROT"	"GO"	"EVIDENCE"	"ONTOLOGY"
[25]	"GOALL"	"EVIDENCEALL"	"ONTOLOGYALL"	"PROBEID"
[29]	"OMIM"	"UCSCKG"		

If we want to extract some sample keys of a particular type, we can use the `keys` method.

```
R> head(keys(hgu95av2.db, keytype="SYMBOL"))
```

[1]	"A1BG"	"A2M"	"A2MP1"	"NAT1"	"NAT2"	"AACP"
-----	--------	-------	---------	--------	--------	--------

And finally, if we have some keys, we can use `select` to extract them. By simply using appropriate argument values with `select` we can specify what keys we want to look up values for (keys), what we want returned back (cols) and the type of keys that we are passing in (keytype)

```
R> #1st get some example keys
R> k <- head(keys(hgu95av2.db, keytype="PROBEID"))
R> # then call select
R> select(hgu95av2.db, keys=k, cols=c("SYMBOL", "GENENAME"), keytype="PROBEID")
```

	PROBEID	SYMBOL		GENENAME
1	1000_at	MAPK3		
2	1001_at	TIE1		
3	1002_f_at	CYP2C19		
4	1003_s_at	CXCR5		
5	1004_at	CXCR5		
6	1005_at	DUSP1		
1				mitogen-activated protein kinase 3
2				tyrosine kinase with immunoglobulin-like and EGF-like domains 1
3				cytochrome P450, family 2, subfamily C, polypeptide 19
4				chemokine (C-X-C motif) receptor 5
5				chemokine (C-X-C motif) receptor 5
6				dual specificity phosphatase 1

And as you can see, when you call the code above, select will try to return a data.frame with all the things you asked for matched up to each other.

0.3 OrgDb objects and the select method

An organism level package (an ‘org’ package) uses a central gene identifier (e.g. Entrez Gene id) and contains mappings between this identifier and other kinds of identifiers (e.g. GenBank or Uniprot accession number, RefSeq id, etc.). The name of an org package is always of the form *org.<Ab>.<id>.db* (e.g. *org.Sc.sgd.db*) where *<Ab>* is a 2-letter abbreviation of the organism (e.g. *Sc* for *Saccharomyces cerevisiae*) and *<id>* is an abbreviation (in lower-case) describing the type of central identifier (e.g. *sgd* for gene identifiers assigned by the Saccharomyces Genome Database, or *eg* for Entrez Gene ids).

Just as the chip packages load a *ChipDb* object, the org packages will load a *OrgDb* object. The following exercise should acquaint you with the use of these methods in the context of an organism package.

Exercise 1

Display the *OrgDb* object for the *org.Hs.eg.db* package.

Use the `cols` method to discover which sorts of annotations can be extracted from it. Is this the same as the result from the `keytypes` method? Use the `keytypes` method to find out.

Finally, use the `keys` method to extract UNIPROT identifiers and then pass those keys in to the `select` method in such a way that you extract the gene symbol and KEGG pathway information for each. Use the help system as needed to learn which values to pass in to `cols` in order to achieve this.

Solution:

```
R> library(org.Hs.eg.db)
R> cols(org.Hs.eg.db)
```

```
[1] "ENTREZID"      "PFAM"          "IPI"           "PROSITE"
[5] "ACCNUM"        "ALIAS"         "CHR"           "CHRLOC"
[9] "CHRLOCEND"     "ENZYME"        "MAP"           "PATH"
[13] "PMID"          "REFSEQ"        "SYMBOL"        "UNIGENE"
[17] "ENSEMBL"       "ENSEMBLPROT"   "ENSEMBLTRANS"  "GENENAME"
[21] "UNIPROT"       "GO"            "EVIDENCE"      "ONTOLOGY"
[25] "GOALL"         "EVIDENCEALL"   "ONTOLOGYALL"   "OMIM"
[29] "UCSCKG"
```

```
R> help("SYMBOL") ## for explanation of these cols and keytypes values
```

```
R> keytypes(org.Hs.eg.db)
```

```
[1] "ENTREZID"      "PFAM"          "IPI"           "PROSITE"
[5] "ACCNUM"        "ALIAS"         "CHR"           "CHRLOC"
[9] "CHRLOCEND"     "ENZYME"        "MAP"           "PATH"
[13] "PMID"          "REFSEQ"        "SYMBOL"        "UNIGENE"
[17] "ENSEMBL"       "ENSEMBLPROT"   "ENSEMBLTRANS"  "GENENAME"
[21] "UNIPROT"       "GO"            "EVIDENCE"      "ONTOLOGY"
[25] "GOALL"         "EVIDENCEALL"   "ONTOLOGYALL"   "OMIM"
[29] "UCSCKG"
```

```
R> uniKeys <- head(keys(org.Hs.eg.db, keytype="UNIPROT"))
R> cols <- c("SYMBOL", "PATH")
R> select(org.Hs.eg.db, keys=uniKeys, cols=cols, keytype="UNIPROT")
```

```
  UNIPROT SYMBOL  PATH
1  P04217  A1BG  <NA>
2  P01023  A2M  04610
```

```

3  F5H5R8  NAT1 00232
4  F5H5R8  NAT1 00983
5  F5H5R8  NAT1 01100
6  P18440  NAT1 00232
7  P18440  NAT1 00983
8  P18440  NAT1 01100
9  Q400J6  NAT1 00232
10 Q400J6  NAT1 00983
11 Q400J6  NAT1 01100
12 A4Z6T7  NAT2 00232
13 A4Z6T7  NAT2 00983
14 A4Z6T7  NAT2 01100

```

So how could you use `select` to annotate your results? This next exercise should help you to understand how that should generally work.

Exercise 2

Please run the following code snippet (which will load a fake data result that I have provided for the purposes of illustration):

```

R> load(system.file("extdata", "resultTable.Rda", package="AnnotationDbi"))
R> head(resultTable)

```

	logConc	logFC	LR.statistic	PValue	FDR
100418920	-9.639471	-4.679498	378.0732	3.269307e-84	2.613484e-80
100419779	-10.638865	-4.264830	291.1028	2.859424e-65	1.142912e-61
100271867	-11.448981	-4.009603	222.3653	2.757135e-50	7.346846e-47
100287169	-11.026699	-3.486593	206.7771	6.934967e-47	1.385953e-43
100287735	-11.036862	3.064980	204.1235	2.630432e-46	4.205535e-43
100421986	-12.276297	-4.695736	190.5368	2.427556e-43	3.234314e-40

The rownames of this table happen to provide entrez gene identifiers for each row (for human). Find the gene symbol and gene name for each of the rows in `resultTable` and then use the `merge` method to attach those annotations to it.

Solution:

```

R> annots <- select(org.Hs.eg.db, keys=rownames(resultTable),
                    cols=c("SYMBOL", "GENENAME"), keytype="ENTREZID")
R> resultTable <- merge(resultTable, annots, by.x=0, by.y="ENTREZID")
R> head(resultTable)

```


	Row.names	logConc	logFC	LR.statistic	PValue	FDR
1	100127888	-10.57050	2.758937	182.8937	1.131473e-41	1.130624e-38
2	100131223	-12.37808	-4.654318	179.2331	7.126423e-41	6.329847e-38
3	100271381	-12.06340	3.511937	188.4824	6.817155e-43	7.785191e-40
4	100271867	-11.44898	-4.009603	222.3653	2.757135e-50	7.346846e-47
5	100287169	-11.02670	-3.486593	206.7771	6.934967e-47	1.385953e-43
6	100287735	-11.03686	3.064980	204.1235	2.630432e-46	4.205535e-43

	SYMBOL	GENENAME
1	LOC100127888	uncharacterized LOC100127888
2	LOC100131223	ADP-ribosylation factor-like 8B pseudogene
3	RPS28P8	ribosomal protein S28 pseudogene 8
4	MPVQTL1	Mean platelet volume QTL1
5	LOC100287169	ubiquitin-conjugating enzyme E2 variant 1-like
6	TTYT13B	testis-specific transcript, Y-linked 13B

0.4 Using select with GO.db

When you load the GO.db package, a *GODb* object is also loaded. This allows you to use the `cols`, `keys`, `keytypes` and `select` methods on the contents of the GO ontology. So if for example, you had a few GO IDs and wanted to know more about it, you could do it like this:

```
R> library(GO.db)
R> GOIDs <- c("GO:0042254", "GO:0044183")
R> select(GO.db, keys=GOIDs, cols="DEFINITION", keytype="GOID")
```

```
GOID
1 GO:0042254
2 GO:0044183
```

```
1 A cellular process that results in the biosynthesis of constituent macromolecu
2 Interacting selectively and non-covalently with any protein or protein complex (a c
```

0.5 Using select with TranscriptDb packages

A *TranscriptDb* package (a 'TxDb' package) connects a set of genomic coordinates to various transcript oriented features. The package can also contain Identifiers to features such as genes and transcripts, and the internal schema describes the relationships between these different elements. All

TranscriptDb containing packages follow a specific naming scheme that tells where the data came from as well as which build of the genome it comes from.

Exercise 3

Display the *TranscriptDb* object for the *TxDb.Hsapiens.UCSC.hg19.knownGene* package.

As before, use the *cols* and *keytypes* methods to discover which sorts of annotations can be extracted from it.

Use the *keys* method to extract just a few gene identifiers and then pass those keys in to the *select* method in such a way that you extract the transcript ids and transcript starts for each.

Solution:

```
R> library(TxDb.Hsapiens.UCSC.hg19.knownGene)
R> txdb <- TxDb.Hsapiens.UCSC.hg19.knownGene
R> txdb
```

TranscriptDb object:

```
| Db type: TranscriptDb
| Supporting package: GenomicFeatures
| Data source: UCSC
| Genome: hg19
| Organism: Homo sapiens
| UCSC Table: knownGene
| Resource URL: http://genome.ucsc.edu/
| Type of Gene ID: Entrez Gene ID
| Full dataset: yes
| miRBase build ID: GRCh37
| transcript_nrow: 80922
| exon_nrow: 286852
| cds_nrow: 235842
| Db created by: GenomicFeatures package from Bioconductor
| Creation time: 2013-03-08 09:43:09 -0800 (Fri, 08 Mar 2013)
| GenomicFeatures version at creation time: 1.11.14
| RSQLite version at creation time: 0.11.2
| DBSCHEMAVERSION: 1.0
```

```
R> cols(txdb)
```

```

[1] "CDSID"      "CDSNAME"    "CDSCHROM"   "CDSSTRAND"  "CDSSTART"
[6] "CDSEND"     "EXONID"     "EXONNAME"   "EXONCHROM"  "EXONSTRAND"
[11] "EXONSTART"  "EXONEND"    "GENEID"     "TXID"       "EXONRANK"
[16] "TXNAME"     "TXCHROM"    "TXSTRAND"   "TXSTART"    "TXEND"

```

```
R> keytypes(txdb)
```

```

[1] "GENEID"    "TXID"      "TXNAME"    "EXONID"    "EXONNAME"  "CDSID"
[7] "CDSNAME"

```

```
R> keys <- head(keys(txdb, keytype="GENEID"))
```

```
R> cols <- c("TXID", "TXSTART")
```

```
R> select(txdb, keys=keys, cols=cols, keytype="GENEID")
```

	GENEID	TXID	TXSTART
1	1	68796	58858172
2	1	68797	58859832
3	10	31203	18248755
4	100	70442	43248163
5	1000	63795	25530930
6	1000	63796	25530930
7	10000	7702	243651535
8	10000	7703	243663021
9	10000	7704	243663021
10	100008586	74057	49217763

```
R>
```

As is widely known, in addition to providing access via the `select` method, *TranscriptDb* objects also provide access via the more familiar `transcripts`, `exons`, `cds`, `transcriptsBy`, `exonsBy` and `cdsBy` methods. For those who do not yet know about these other methods, more can be learned by seeing the vignette called: *Making and Utilizing TranscriptDb Objects* in the *GenomicFeatures* package.

The version number of R and packages loaded for generating the vignette were:

```

R version 3.0.0 (2013-04-03)
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=en_US.UTF-8  LC_MESSAGES=en_US.UTF-8
[7] LC_PAPER=C               LC_NAME=C
[9] LC_ADDRESS=C             LC_TELEPHONE=C
[11] LC_MEASUREMENT=en_US.UTF-8 LC_IDENTIFICATION=C
```

attached base packages:

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

other attached packages:

```
[1] TxDb.Hsapiens.UCSC.hg19.knownGene_2.9.0
[2] GenomicFeatures_1.12.0
[3] GenomicRanges_1.12.1
[4] IRanges_1.18.0
[5] GO.db_2.9.0
[6] hgu95av2.db_2.9.0
[7] AnnotationForge_1.2.0
[8] org.Hs.eg.db_2.9.0
[9] RSQLite_0.11.2
[10] DBI_0.2-5
[11] AnnotationDbi_1.22.1
[12] Biobase_2.20.0
[13] BiocGenerics_0.6.0
```

loaded via a namespace (and not attached):

```
[1] BSgenome_1.28.0 Biostrings_2.28.0 RCurl_1.95-4.1
[4] Rsamtools_1.12.0 XML_3.96-1.1 biomaRt_2.16.0
[7] bitops_1.0-5 rtracklayer_1.20.0 stats4_3.0.0
[10] tools_3.0.0 zlibbioc_1.6.0
```