

Creating a New Annotation Package using SQLForge

Marc Carlson, Herve Pages, Nianhua Li

July 4, 2013

1 Introduction

The *AnnotationForge* package provides a series of functions that can be used to build *AnnotationDbi* packages for supported organisms. This collection of functions is called SQLForge.

In order to use SQLForge you really only need to have one kind of information and that is a list of paired IDs. These IDs are to be stored in a tab delimited file that is formatted in the same way that they used to be for the older AnnBuilder package. For those who are unfamiliar with the AnnBuilder package, this just means that there are two columns separated by a tab where the column on the left contains probe or probeset identifiers and the column on the right contains some sort of widely accepted gene accession. This file should NOT contain a header. SQLForge will then use these IDs along with it's own support databases to make an *AnnotationDbi* package for you. Here is how these IDs should look if you were to read them into R:

```
> library(AnnotationForge)
> read.table(system.file("extdata", "hcg110_ID",
                        package="AnnotationDbi"),
             sep = "\t", header = FALSE, as.is = TRUE)[1:5,]
```

	V1	V2
1	1000_at	X60188
2	1001_at	X60957
3	1002_f_at	X65962
4	1003_s_at	X68149
5	1004_at	X68149

In the example above, Genbank IDs are demonstrated. But it is also possible to use entrez gene IDs, refseq IDs or unigene accessions as the gene identifiers. If refseq IDs are used, it is preferable to strip off the version extensions that can sometimes be added on by some vendors. The version extensions are digits that are sometimes tacked onto the end of a refseq ID and separated from the accession by a dot. As an example consider "NM_000193.2" . The "NM_000193" portion would be the actual accession number and the ".2" would be the version number. These version numbers are not used by these databases and their presence in your input can cause less than desirable results.

Alternatively, if you have an annotation file for an Affymetrix chip, you can use a parameter called `affy` that will automatically parse such a file and produce a similar mapping from that. It is important to understand however that despite that rather rich contents of an Affymetrix annotation file, almost none of these data are used in the making of an annotation package with SQLForge. Instead, the relevant IDs are stripped out, and then passed along to SQLForge as if you had created a file like is seen above. The option here to use such a file is offered purely as a convenience because the platform is so popular.

If you have additional information about your probes in the form of other kinds of supported gene IDs, you can pass these in as well by using the `otherSrc` parameter. These IDs must be formatted using the same two column format as described above, and if there are multiple source files, then you can pass them in as a list of strings that correspond to the file paths for these files.

Once you have your IDs ready, SQLForge will read them in, and use the gene IDs to compare to an intermediate database. The data from this database is what is used to make the specialized database that is placed inside of an annotation package.

At the present time, it is possible to make annotation packages for the most common model organisms. For each of these organisms another support package will be maintained and updated biannually which will include all the basic data gathered for this organism from sources such as NCBI, GO, KEGG and Flybase etc. These support packages will each be named after the organism they are intended for and will each include a large sqlite database with all the supporting information for that organism. Please note that support databases are not necessary unless you intend to actually make a new annotation package for one of the supported organisms. In the case where you want to make annotation packages, the support databases are only required for the organism in question. When SQLForge makes a new

database, it uses the information supplied by the support database as the data source to make the annotation package. So the relevant support packages need to be updated to the latest version in order to guarantee that the annotation packages you produce will be made with information from the last biannual update. These support packages are not meant to be annotation packages themselves and they come with no schema of their own. Instead these are merely a way to distribute the data to those who want to make custom annotation packages.

To check if your organism is supported simply look in the metadata packages repository on the bioconductor website for a .db0 package. Only special organism base packages will end with the .db0 extension. If you find a package that is named after the organism you are interested in, then your organism is supported, and you can use that database to make custom packages. To list all the supported organism .db0 packages directly from R you can use `available.db0pkgs()`.

```
> available.db0pkgs()

[1] "anopheles.db0" "arabidopsis.db0" "bovine.db0"
[4] "canine.db0"     "chicken.db0"     "chimp.db0"
[7] "ecoliK12.db0"   "ecoliSakai.db0" "fly.db0"
[10] "human.db0"      "malaria.db0"     "mouse.db0"
[13] "pig.db0"        "rat.db0"         "rhesus.db0"
[16] "worm.db0"       "xenopus.db0"     "yeast.db0"
[19] "zebrafish.db0"
```

2 How to use SQLForge

Once you know the name of the package that you need, you can get the latest .db0 package for your organism by using `biocLite()` like this.

```
> source("http://bioconductor.org/biocLite.R")
> biocLite("human.db0")
```

Even if you have installed the .db0 package before, it's a good idea to run this again to be sure that you have the latest organism package.

Once you have the appropriate .db0 package, you may also need to install the appropriate `org.*` package. Most users probably have this already, because using any chip packages for this organism will require it. But if you don't have it yet, you can find a list of these packages on the web site here:

```
http://www.bioconductor.org/packages/release/BiocViews.html#
---OrgDb
```

Once you know the name of the package you need, you can install it like before.

```
> biocLite("org.Hs.eg.db")
```

All the supported organisms with a .db0 package, will also have a matching org.* package. The org.* packages are named like this: org.<species abbreviation>.<source abbreviation>.db. The source indicates where the data originates from and also which central ID the data all connects to, while the species abbreviation indicates which organism the data is for. The web site spells out the species a bit more clearly by listing the titles for each package.

Since each organism will have different kinds of data available, the schemas that will be needed for each organism will also change. SQLForge provides support functions for each of the model organisms that will create a sqlite database that complies with a specified database schema. The following helper function will list supported Schemas.

```
> available.dbschemas()

[1] "ARABIDOPSISCHIP_DB" "BOVINECHIP_DB"
[3] "BOVINE_DB"          "CANINECHIP_DB"
[5] "CANINE_DB"          "CHICKENCHIP_DB"
[7] "CHICKEN_DB"         "ECOLICHIP_DB"
[9] "ECOLI_DB"           "FLYCHIP_DB"
[11] "FLY_DB"             "GO_DB"
[13] "HUMANCHIP_DB"       "HUMANCROSSCHIP_DB"
[15] "HUMAN_DB"           "INPARANOIDDROME_DB"
[17] "INPARANOIDHOMSA_DB" "INPARANOIDMUSMU_DB"
[19] "INPARANOIDRATNO_DB" "INPARANOIDSACCE_DB"
[21] "KEGG_DB"            "MALARIA_DB"
[23] "MOUSECHIP_DB"       "MOUSE_DB"
[25] "PFAM_DB"            "PIGCHIP_DB"
[27] "PIG_DB"             "RATCHIP_DB"
[29] "RAT_DB"             "WORMCHIP_DB"
[31] "WORM_DB"            "YEASTCHIP_DB"
[33] "YEAST_DB"           "ZEBRAFISHCHIP_DB"
[35] "ZEBRAFISH_DB"
```

The following shows an example of how to make a chip package:

```
> hcg110_IDs = system.file("extdata",  
                           "hcg110_ID",  
                           package="AnnotationDbi")  
  
> tmpout = tempdir()  
> makeDBPackage("HUMANCHIP_DB",  
               affy=FALSE,  
               prefix="hcg110",  
               fileName=hcg110_IDs,  
               baseMapType="gb",  
               outputDir = tmpout,  
               version="1.0.0",  
               manufacturer = "Affymetrix",  
               chipName = "Human Cancer G110 Array",  
               manufacturerUrl = "http://www.affymetrix.com")
```

Wrapper functions are provided for making all of the different kinds of chip based package types that are presently defined. These are named after the schemas that they correspond to. So for example `makeHUMANCHIP_DB()` corresponds to the HUMANCHIP_DB schema, and is used to produce chip based annotation packages of that type.

2.1 Installing your custom package

To install your package in Unix simply use R CMD INSTALL <package-Name> at the command line. But if you are on Windows or Mac, you may have to instead use `install.packages` from within R. This will work because this kind of simple annotation package does not contain any code that has to be compiled. So you can simply call `install.packages` and set the `repos` parameter to `NULL` and the `type` parameter to "source". The final R command will look something like this:

```
> install.packages("packageNameAndPath", repos=NULL, type="source")
```

Of course, you still have to type the path to your source directory correctly as the 1st argument. It is recommended that you use the autocomplete feature in R as you enter it so that you get the path specified correctly.

3 For Advanced users: How to add extra data into your packages

Sometimes you may find that you want to add extra supplementary data into the database for the package that you just created. In these cases, you will have to begin by using the SQL to add more data into the database. Before you can do that however, you will have to change the permissions on the sqlite database. The database will always be in the `inst/extdata` directory of your package source after you run `SQLForge`. Once you can edit your database, you will have to create a new table, and populate that table with new information using SQL statements. One good way to do this would be to use the *RSQLite* interface that is introduced in portions of the *AnnotationDbi* vignette. For a more thorough treatment of the *RSQLite* package, please see the vignette for that package at CRAN. Once you are finished editing the database with SQL, be sure to change the database file back to being a read only file.

However, adding the content to the database is only the 1st part of what has to be done. In order for the data to be exposed to the R layer as a mapping, you will have to also create and document a mapping object. To do this step we have added a simple utility function to *AnnotationDbi* that allows you to make a simple *Bimap* from a single table. The following example will make an additional mapping between the gene names and the gene symbols found in the `gene_info` table for the package *hgu95av2.db*. For this particular example, no additional SQL has to be inserted 1st into the database since it is just adding a mapping onto data that already exists in the database (but is just not normally exposed as a mapping).

```
> library(hgu95av2.db)
> hgu95av2NAMESYMBOL <- createSimpleBimap("gene_info",
                                         "gene_name",
                                         "symbol",
                                         hgu95av2.db::datacache,
                                         "NAMESYMBOL",
                                         "hgu95av2.db")

> ##What is the mapping we just made?
> hgu95av2NAMESYMBOL

NAMESYMBOL map for hgu95av2.db (object of class "AnnDbBimap")

> ##Display the 1st 4 relationships in this new mapping
> as.list(hgu95av2NAMESYMBOL)[1:4]
```

```
$`alpha-1-B glycoprotein`
```

```
[1] "A1BG"
```

```
$`alpha-2-macroglobulin`
```

```
[1] "A2M"
```

```
$`alpha-2-macroglobulin pseudogene 1`
```

```
[1] "A2MP1"
```

```
$`N-acetyltransferase 1 (arylamine N-acetyltransferase)`
```

```
[1] "NAT1"
```

If instead of creating a mapping on an existing example, you wanted to add a new mapping to your customized annotation package, you would need to call this function from `zzz.R` in your modified annotation package (and also expose it in the namespace). You will then want to be sure that your updated database has replaced the one in the `inst/extdata` directory that was originally generated by `SQLForge`. And finally, you will need to also put a man page into your package so that users will know how to make use of this new mapping.

4 Session Information

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

```
R version 3.0.1 (2013-05-16)
```

```
Platform: i386-w64-mingw32/i386 (32-bit)
```

```
locale:
```

```
[1] LC_COLLATE=C
```

```
[2] LC_CTYPE=English_United States.1252
```

```
[3] LC_MONETARY=English_United States.1252
```

```
[4] LC_NUMERIC=C
```

```
[5] LC_TIME=English_United States.1252
```

```
attached base packages:
```

```
[1] parallel stats graphics grDevices utils
```

```
[6] datasets methods base
```

other attached packages:

- [1] hgu95av2.db_2.9.0
- [2] AnnotationForge_1.2.2
- [3] hom.Hs.inp.db_2.9.1
- [4] Biostrings_2.28.0
- [5] XML_3.98-1.1
- [6] Homo.sapiens_1.1.1
- [7] TxDb.Hsapiens.UCSC.hg19.knownGene_2.9.2
- [8] org.Hs.eg.db_2.9.0
- [9] GO.db_2.9.0
- [10] RSQLite_0.11.4
- [11] DBI_0.2-7
- [12] OrganismDbi_1.2.0
- [13] GenomicFeatures_1.12.2
- [14] GenomicRanges_1.12.4
- [15] IRanges_1.18.1
- [16] AnnotationDbi_1.22.6
- [17] Biobase_2.20.1
- [18] BiocGenerics_0.6.0

loaded via a namespace (and not attached):

- [1] BSgenome_1.28.0 RBGL_1.36.2
- [3] RCurl_1.95-4.1 Rsamtools_1.12.3
- [5] biomaRt_2.16.0 bitops_1.0-5
- [7] graph_1.38.2 rtracklayer_1.20.2
- [9] stats4_3.0.1 tools_3.0.1
- [11] zlibbioc_1.6.0