

porcine.db

April 2, 2013

porcineACCNUM *Map Manufacturer identifiers to Accession Numbers*

Description

porcineACCNUM is an R object that contains mappings between a manufacturer's identifiers and manufacturers accessions.

Details

For chip packages such as this, the ACCNUM mapping comes directly from the manufacturer. This is different from other mappings which are mapped onto the probes via an Entrez Gene identifier.

Each manufacturer identifier maps to a vector containing a GenBank accession number.

Mappings were based on data provided by: Entrez Gene <ftp://ftp.ncbi.nlm.nih.gov/gene/DATA> With a date stamp from the source of: 2012-Sep4

Examples

```
x <- porcineACCNUM
# Get the probe identifiers that are mapped to an ACCNUM
mapped_probes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_probes])
if(length(xx) > 0) {
  # Get the ACCNUM for the first five probes
  xx[1:5]
  # Get the first one
  xx[[1]]
}
```

porcineALIAS2PROBE *Map between Common Gene Symbol Identifiers and Manufacturer Identifiers*

Description

porcineALIAS is an R object that provides mappings between common gene symbol identifiers and manufacturer identifiers.

Details

Each gene symbol is mapped to a named vector of manufacturer identifiers. The name represents the gene symbol and the vector contains all manufacturer identifiers that are found for that symbol. An NA is reported for any gene symbol that cannot be mapped to any manufacturer identifiers.

This mapping includes ALL gene symbols including those which are already listed in the SYMBOL map. The SYMBOL map is meant to only list official gene symbols, while the ALIAS maps are meant to store all used symbols.

Mappings were based on data provided by: Entrez Gene <ftp://ftp.ncbi.nlm.nih.gov/gene/DATA> With a date stamp from the source of: 2012-Sep4

Examples

```
# Convert the object to a list
xx <- as.list(porcineALIAS2PROBE)
if(length(xx) > 0){
  # Get the probe identifiers for the first two aliases
  xx[1:2]
  # Get the first one
  xx[[1]]
}
```

porcine.db

Bioconductor annotation data package

Description

Welcome to the porcine.db annotation Package. The purpose of this package is to provide detailed information about the porcine platform. This package is updated biannually.

You can learn what objects this package supports with the following command:

```
ls("package:porcine.db")
```

Each of these objects has their own manual page detailing where relevant data was obtained along with some examples of how to use it.

Examples

```
ls("package:porcine.db")
```

porcineCHR

Map Manufacturer IDs to Chromosomes

Description

porcineCHR is an R object that provides mappings between a manufacturer identifier and the chromosome that contains the gene of interest.

Details

Each manufacturer identifier maps to a vector of chromosomes. Due to inconsistencies that may exist at the time the object was built, the vector may contain more than one chromosome (e.g., the identifier may map to more than one chromosome). If the chromosomal location is unknown, the vector will contain an NA.

Mappings were based on data provided by: Entrez Gene <ftp://ftp.ncbi.nlm.nih.gov/gene/DATA> With a date stamp from the source of: 2012-Sep4

Examples

```
x <- porcineCHR
# Get the probe identifiers that are mapped to a chromosome
mapped_probes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_probes])
if(length(xx) > 0) {
  # Get the CHR for the first five probes
  xx[1:5]
  # Get the first one
  xx[[1]]
}
```

porcine_dbconn

Collect information about the package annotation DB

Description

Some convenience functions for getting a connection object to (or collecting information about) the package annotation DB.

Usage

```
porcine_dbconn()
porcine_dbfile()
porcine_dbschema(file="", show.indices=FALSE)
porcine_dbInfo()
```

Arguments

file	A connection, or a character string naming the file to print to (see the file argument of the cat function for the details).
show.indices	The CREATE INDEX statements are not shown by default. Use show.indices=TRUE to get them.

Details

porcine_dbconn returns a connection object to the package annotation DB. **IMPORTANT:** Don't call `dbDisconnect` on the connection object returned by `porcine_dbconn` or you will break all the `AnnDbObj` objects defined in this package!

porcine_dbfile returns the path (character string) to the package annotation DB (this is an SQLite file).

porcine_dbschema prints the schema definition of the package annotation DB.

porcine_dbInfo prints other information about the package annotation DB.

Value

porcine_dbconn: a `DBConnection` object representing an open connection to the package annotation DB.

porcine_dbfile: a character string with the path to the package annotation DB.

porcine_dbschema: none (invisible NULL).

porcine_dbInfo: none (invisible NULL).

See Also

[dbGetQuery](#), [dbConnect](#), [dbconn](#), [dbfile](#), [dbschema](#), [dbInfo](#)

Examples

```
## Count the number of rows in the "probes" table:
dbGetQuery(porcine_dbconn(), "SELECT COUNT(*) FROM probes")

## The connection object returned by porcine_dbconn() was
## created with:
dbConnect(SQLite(), dbname=porcine_dbfile(), cache_size=64000,
synchronous=0)

porcine_dbschema()

porcine_dbInfo()
```

porcineENTREZID *Map between Manufacturer Identifiers and Entrez Gene*

Description

porcineENTREZID is an R object that provides mappings between manufacturer identifiers and Entrez Gene identifiers.

Details

Each manufacturer identifier is mapped to a vector of Entrez Gene identifiers. An NA is assigned to those manufacturer identifiers that can not be mapped to an Entrez Gene identifier at this time.

If a given manufacturer identifier can be mapped to different Entrez Gene identifiers from various sources, we attempt to select the common identifiers. If a consensus cannot be determined, we select the smallest identifier.

Mappings were based on data provided by: Entrez Gene <ftp://ftp.ncbi.nlm.nih.gov/gene/DATA> With a date stamp from the source of: 2012-Sep4

References

<http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?db=gene>

Examples

```
x <- porcineENTREZID
# Get the probe identifiers that are mapped to an ENTREZ Gene ID
mapped_probes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_probes])
if(length(xx) > 0) {
  # Get the ENTREZID for the first five probes
  xx[1:5]
  # Get the first one
  xx[[1]]
}
```

porcineENZYME

Maps between Manufacturer IDs and Enzyme Commission (EC) Numbers

Description

porcineENZYME is an R object that provides mappings between manufacturer identifiers and EC numbers. porcineENZYME2PROBE is an R object that maps Enzyme Commission (EC) numbers to manufacturer identifiers.

Details

When the porcineENZYME mapping is viewed as a list, each manufacturer identifier maps to a named vector containing the EC number that corresponds to the enzyme produced by that gene. The names correspond to the manufacturer identifiers. If this information is unknown, the vector will contain an NA.

For the porcineENZYME2PROBE, each EC number maps to a named vector containing all of the manufacturer identifiers that correspond to the gene that produces that enzyme. The name of the vector corresponds to the EC number.

Enzyme Commission numbers are assigned by the Nomenclature Committee of the International Union of Biochemistry and Molecular Biology <http://www.chem.qmw.ac.uk/iubmb/enzyme/> to allow enzymes to be identified.

An Enzyme Commission number is of the format EC x.y.z.w, where x, y, z, and w are numeric numbers. In porcineENZYME2PROBE, EC is dropped from the Enzyme Commission numbers.

Enzyme Commission numbers have corresponding names that describe the functions of enzymes in such a way that EC x is a more general description than EC x.y that in turn is a more general description than EC x.y.z. The top level EC numbers and names are listed below:

EC 1 oxidoreductases

EC 2 transferases

EC 3 hydrolases

EC 4 lyases

EC 5 isomerases

EC 6 ligases

The EC name for a given EC number can be viewed at <http://www.chem.qmul.ac.uk/iupac/jcbrn/index.html#6>

Mappings between probe identifiers and enzyme identifiers were obtained using files provided by: KEGG GENOME <ftp://ftp.genome.jp/pub/kegg/genomes> With a date stamp from the source of: 2011-Mar15

References

<ftp://ftp.genome.ad.jp/pub/kegg/pathways>

Examples

```
x <- porcineENZYME
# Get the probe identifiers that are mapped to an EC number
mapped_probes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_probes])
if(length(xx) > 0) {
  # Get the ENZYME for the first five probes
  xx[1:5]
  # Get the first one
  xx[[1]]
}

# Now convert porcineENZYME2PROBE to a list to see inside
xx <- as.list(porcineENZYME2PROBE)
if(length(xx) > 0){
  # Get the probe identifiers for the first five enzyme
  #commission numbers
  xx[1:5]
  # Get the first one
  xx[[1]]
}
```

porcineGENENAME *Map between Manufacturer IDs and Genes*

Description

porcineGENENAME is an R object that maps manufacturer identifiers to the corresponding gene name.

Details

Each manufacturer identifier maps to a named vector containing the gene name. The vector name corresponds to the manufacturer identifier. If the gene name is unknown, the vector will contain an NA.

Gene names currently include both the official (validated by a nomenclature committee) and preferred names (interim selected for display) for genes. Efforts are being made to differentiate the two by adding a name to the vector.

Mappings were based on data provided by: Entrez Gene <ftp://ftp.ncbi.nlm.nih.gov/gene/DATA> With a date stamp from the source of: 2012-Sep4

Examples

```
x <- porcineGENENAME
# Get the probe identifiers that are mapped to a gene name
mapped_probes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_probes])
if(length(xx) > 0) {
  # Get the GENENAME for the first five probes
  xx[1:5]
  # Get the first one
  xx[[1]]
}
```

porcineGO

Maps between manufacturer IDs and Gene Ontology (GO) IDs

Description

porcineGO is an R object that provides mappings between manufacturer identifiers and the GO identifiers that they are directly associated with. This mapping and its reverse mapping (porcineGO2PROBE) do NOT associate the child terms from the GO ontology with the gene. Only the directly evidenced terms are represented here.

porcineGO2ALLPROBES is an R object that provides mappings between a given GO identifier and all of the manufacturer identifiers annotated at that GO term OR TO ONE OF IT'S CHILD NODES in the GO ontology. Thus, this mapping is much larger and more inclusive than porcineGO2PROBE.

Details

If porcineGO is cast as a list, each manufacturer identifier is mapped to a list of lists. The names on the outer list are GO identifiers. Each inner list consists of three named elements: GOID, Ontology, and Evidence.

The GOID element matches the GO identifier named in the outer list and is included for convenience when processing the data using 'lapply'.

The Ontology element indicates which of the three Gene Ontology categories this identifier belongs to. The categories are biological process (BP), cellular component (CC), and molecular function (MF).

The Evidence element contains a code indicating what kind of evidence supports the association of the GO identifier to the manufacturer id. The evidence codes in use include:

IMP: inferred from mutant phenotype

IGI: inferred from genetic interaction

IPI: inferred from physical interaction

ISS: inferred from sequence similarity

IDA: inferred from direct assay

IEP: inferred from expression pattern

IEA: inferred from electronic annotation

TAS: traceable author statement

NAS: non-traceable author statement

ND: no biological data available

IC: inferred by curator

If porcineGO2ALLPROBES or porcineGO2PROBE is cast as a list, each GO term maps to a named vector of manufacturer identifiers and evidence codes. A GO identifier may be mapped to the same manufacturer identifier more than once but the evidence code can be different. Mappings between Gene Ontology identifiers and Gene Ontology terms and other information are available in a separate data package named GO.

Whenever any of these mappings are cast as a data.frame, all the results will be output in an appropriate tabular form.

Mappings between manufacturer identifiers and GO information were obtained through their mappings to manufacturer identifiers. NAs are assigned to manufacturer identifiers that can not be mapped to any Gene Ontology information. Mappings between Gene Ontology identifiers and Gene Ontology terms and other information are available in a separate data package named GO.

All mappings were based on data provided by: Gene Ontology <ftp://ftp.geneontology.org/pub/go/godatabase/archive/latest/lite/> With a date stamp from the source of: 20120901

References

<ftp://ftp.ncbi.nlm.nih.gov/gene/DATA/>

See Also

[porcineGO2ALLPROBES](#).

Examples

```
x <- porcineGO
# Get the manufacturer identifiers that are mapped to a GO ID
mapped_genes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_genes])
if(length(xx) > 0) {
  # Try the first one
  got <- xx[[1]]
  got[[1]][["GOID"]]
  got[[1]][["Ontology"]]
  got[[1]][["Evidence"]]
}
# For the reverse map:
# Convert to a list
xx <- as.list(porcineGO2PROBE)
if(length(xx) > 0){
  # Gets the manufacturer ids for the top 2nd and 3rd GO identifiers
  goids <- xx[2:3]
  # Gets the manufacturer ids for the first element of goids
  goids[[1]]
  # Evidence code for the mappings
```



```
    names(goids[[1]])
  }
  # Convert porcineGO2ALLPROBES to a list
  xx <- as.list(porcineGO2ALLPROBES)
  if(length(xx) > 0){
    # Gets the manufacturer identifiers for the top 2nd and 3rd GO identifiers
    goids <- xx[2:3]
    # Gets all the manufacturer identifiers for the first element of goids
    goids[[1]]
    # Evidence code for the mappings
    names(goids[[1]])
  }
```

porcineMAPCOUNTS *Number of mapped keys for the maps in package porcine.db*

Description

porcineMAPCOUNTS provides the "map count" (i.e. the count of mapped keys) for each map in package porcine.db.

Details

This "map count" information is precalculated and stored in the package annotation DB. This allows some quality control and is used by the [checkMAPCOUNTS](#) function defined in AnnotationDbi to compare and validate different methods (like `count.mappedkeys(x)` or `sum(!is.na(as.list(x)))`) for getting the "map count" of a given map.

See Also

[mappedkeys](#), [count.mappedkeys](#), [checkMAPCOUNTS](#)

Examples

```
porcineMAPCOUNTS
mapnames <- names(porcineMAPCOUNTS)
porcineMAPCOUNTS[mapnames[1]]
x <- get(mapnames[1])
sum(!is.na(as.list(x)))
count.mappedkeys(x) # much faster!

## Check the "map count" of all the maps in package porcine.db
checkMAPCOUNTS("porcine.db")
```

porcineORGANISM *The Organism information for porcine*

Description

porcineORGANISM is an R object that contains a single item: a character string that names the organism for which porcine was built. porcineORGPKG is an R object that contains a character vector with the name of the organism package that a chip package depends on for its gene-centric annotation.

Details

Although the package name is suggestive of the organism for which it was built, porcineORGANISM provides a simple way to programmatically extract the organism name. porcineORGPKG provides a simple way to programmatically extract the name of the parent organism package. The parent organism package is a strict dependency for chip packages as this is where the gene centric information is ultimately extracted from. The full package name will always be this string plus the extension ".db". But most programatic acces will not require this extension, so its more convenient to leave it out.

Examples

```
porcineORGANISM
porcineORGPKG
```

porcinePATH *Mappings between probe identifiers and KEGG pathway identifiers*

Description

KEGG (Kyoto Encyclopedia of Genes and Genomes) maintains pathway data for various organisms.

porcinePATH maps probe identifiers to the identifiers used by KEGG for pathways in which the genes represented by the probe identifiers are involved

porcinePATH2PROBE is an R object that provides mappings between KEGG identifiers and manufacturer identifiers.

Details

Each KEGG pathway has a name and identifier. Pathway name for a given pathway identifier can be obtained using the KEGG data package that can either be built using AnnBuilder or downloaded from Bioconductor <http://www.bioconductor.org>.

Graphic presentations of pathways are searchable at url <http://www.genome.ad.jp/kegg/pathway.html> by using pathway identifiers as keys.

Mappings were based on data provided by: KEGG GENOME <ftp://ftp.genome.jp/pub/kegg/genomes>
With a date stamp from the source of: 2011-Mar15

References

<http://www.genome.ad.jp/kegg/>

Examples

```
x <- porcinePATH
# Get the probe identifiers that are mapped to a KEGG pathway ID
mapped_probes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_probes])
if(length(xx) > 0) {
  # Get the PATH for the first five probes
  xx[1:5]
  # Get the first one
  xx[[1]]
}

# Now convert the porcinePATH2PROBE object to a list
xx <- as.list(porcinePATH2PROBE)
if(length(xx) > 0){
  # Get the probe identifiers for the first two pathway identifiers
  xx[1:2]
  # Get the first one
  xx[[1]]
}
```

porcinePMID

Maps between Manufacturer Identifiers and PubMed Identifiers

Description

porcinePMID is an R object that provides mappings between manufacturer identifiers and PubMed identifiers. porcinePMID2PROBE is an R object that provides mappings between PubMed identifiers and manufacturer identifiers.

Details

When porcinePMID is viewed as a list each manufacturer identifier is mapped to a named vector of PubMed identifiers. The name associated with each vector corresponds to the manufacturer identifier. The length of the vector may be one or greater, depending on how many PubMed identifiers a given manufacturer identifier is mapped to. An NA is reported for any manufacturer identifier that cannot be mapped to a PubMed identifier.

When porcinePMID2PROBE is viewed as a list each PubMed identifier is mapped to a named vector of manufacturer identifiers. The name represents the PubMed identifier and the vector contains all manufacturer identifiers that are represented by that PubMed identifier. The length of the vector may be one or longer, depending on how many manufacturer identifiers are mapped to a given PubMed identifier.

Titles, abstracts, and possibly full texts of articles can be obtained from PubMed by providing a valid PubMed identifier. The pubmed function of annotate can also be used for the same purpose.

Mappings were based on data provided by: Entrez Gene <ftp://ftp.ncbi.nlm.nih.gov/gene/DATA> With a date stamp from the source of: 2012-Sep4

References

<http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?db=PubMed>

Examples

```
x <- porcinePMID
# Get the probe identifiers that are mapped to any PubMed ID
mapped_probes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_probes])
if(length(xx) > 0){
  # Get the PubMed identifiers for the first two probe identifiers
  xx[1:2]
  # Get the first one
  xx[[1]]
  if(interactive() && !is.null(xx[[1]]) && !is.na(xx[[1]])
    && require(annotate)){
    # Get article information as XML files
    xmls <- pubmed(xx[[1]], disp = "data")
    # View article information using a browser
    pubmed(xx[[1]], disp = "browser")
  }
}

# Now convert the reverse map object porcinePMID2PROBE to a list
xx <- as.list(porcinePMID2PROBE)
if(length(xx) > 0){
  # Get the probe identifiers for the first two PubMed identifiers
  xx[1:2]
  # Get the first one
  xx[[1]]
  if(interactive() && require(annotate)){
    # Get article information as XML files for a PubMed id
    xmls <- pubmed(names(xx)[1], disp = "data")
    # View article information using a browser
    pubmed(names(xx)[1], disp = "browser")
  }
}
```

porcineREFSEQ

Map between Manufacturer Identifiers and RefSeq Identifiers

Description

porcineREFSEQ is an R object that provides mappings between manufacturer identifiers and RefSeq identifiers.

Details

Each manufacturer identifier is mapped to a named vector of RefSeq identifiers. The name represents the manufacturer identifier and the vector contains all RefSeq identifiers that can be mapped to that manufacturer identifier. The length of the vector may be one or greater, depending on how many RefSeq identifiers a given manufacturer identifier can be mapped to. An NA is reported for any manufacturer identifier that cannot be mapped to a RefSeq identifier at this time.

References

<http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?db=gene>

Examples

```
x <- porcineSYMBOL
# Get the probe identifiers that are mapped to a gene symbol
mapped_probes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_probes])
if(length(xx) > 0) {
  # Get the SYMBOL for the first five probes
  xx[1:5]
  # Get the first one
  xx[[1]]
}
```

porcineUNIGENE

Map between Manufacturer Identifiers and UniGene cluster identifiers

Description

porcineUNIGENE is an R object that provides mappings between manufacturer identifiers and UniGene identifiers.

Details

Each manufacturer identifier is mapped to a UniGene identifier. An NA is reported if the manufacturer identifier cannot be mapped to UniGene at this time.

A UniGene identifier represents a cluster of sequences of a gene. Using UniGene identifiers one can query the UniGene database for information about the sequences or the Entrez Gene database for information about the genes.

Mappings were based on data provided by: Entrez Gene <ftp://ftp.ncbi.nlm.nih.gov/gene/DATA> With a date stamp from the source of: 2012-Sep4

References

<http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?db=gene>

Examples

```
x <- porcineUNIGENE
# Get the probe identifiers that are mapped to an UNIGENE ID
mapped_probes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_probes])
if(length(xx) > 0) {
  # Get the UNIGENE for the first five probes
  xx[1:5]
  # Get the first one
  xx[[1]]
}
```

porcineUNIPROT

Map Uniprot accession numbers with Entrez Gene identifiers

Description

porcineUNIPROT is an R object that contains mappings between the manufacturer identifiers and Uniprot accession numbers.

Details

This object is a simple mapping of manufacturer identifiers to Uniprot Accessions.

Mappings were based on data provided by NCBI (link above) with an exception for fly, which required retrieving the data from ensembl <http://www.ensembl.org/biomart/martview/>

Examples

```
x <- porcineUNIPROT
# Get the entrez gene IDs that are mapped to an Uniprot ID
mapped_genes <- mappedkeys(x)
# Convert to a list
xx <- as.list(x[mapped_genes])
if(length(xx) > 0) {
  # Get the Uniprot IDs for the first five genes
  xx[1:5]
  # Get the first one
  xx[[1]]
}
```

Index

*Topic **datasets**

- porcine.db, 2
- porcine_dbconn, 3
- porcineACCNUM, 1
- porcineALIAS2PROBE, 2
- porcineCHR, 3
- porcineENTREZID, 4
- porcineENZYME, 5
- porcineGENENAME, 6
- porcineGO, 7
- porcineMAPCOUNTS, 9
- porcineORGANISM, 10
- porcinePATH, 10
- porcinePMID, 11
- porcineREFSEQ, 12
- porcineSYMBOL, 13
- porcineUNIGENE, 14
- porcineUNIPROT, 15

*Topic **utilities**

- porcine_dbconn, 3

AnnDbObj, 4

cat, 3

checkMAPCOUNTS, 9

count.mappedkeys, 9

dbconn, 4

dbConnect, 4

dbDisconnect, 4

dbfile, 4

dbGetQuery, 4

dbInfo, 4

dbschema, 4

mappedkeys, 9

porcine (porcine.db), 2

porcine.db, 2

porcine_dbconn, 3

porcine_dbfile (porcine_dbconn), 3

porcine_dbInfo (porcine_dbconn), 3

porcine_dbschema (porcine_dbconn), 3

porcineACCNUM, 1

porcineALIAS2PROBE, 2

porcineCHR, 3

porcineENTREZID, 4

porcineENZYME, 5

porcineENZYME2PROBE
(porcineENZYME), 5

porcineGENENAME, 6

porcineGO, 7

porcineGO2ALLPROBES, 8

porcineGO2ALLPROBES (porcineGO), 7

porcineGO2PROBE (porcineGO), 7

porcineLOCUSID (porcineENTREZID), 4

porcineMAPCOUNTS, 9

porcineORGANISM, 10

porcineORGPKG (porcineORGANISM), 10

porcinePATH, 10

porcinePATH2PROBE (porcinePATH), 10

porcinePMID, 11

porcinePMID2PROBE (porcinePMID), 11

porcineREFSEQ, 12

porcineSYMBOL, 13

porcineUNIGENE, 14

porcineUNIPROT, 15

porcineUNIPROT2PROBE
(porcineUNIPROT), 15