

HowTo: Querying online Data

Jeff Gentry and Robert Gentleman

November 29, 2006

1 Overview

This article demonstrates how you can make use of the tools that have been provided for on-line querying of data resources. These tools rely on others (such as the NLM and NCBI) providing and documenting appropriate web interfaces. The tools described here allow you to either retrieve the data (usually in XML) or have it rendered in a browser on the local machine. To do this you will need the *Biobase*, *XML*, and *annotate* packages. The functionality in this article was first described in (Gentleman and Gentry, 2002), although some enhancements have been made since the writing of that article.

Assembling and using meta-data annotation is a non-trivial task. In the Bioconductor Project we have developed tools to support two different methods of accessing meta-data. One is based on obtaining data from a variety of sources, of curating it and packaging it in a form that is suitable for analysing microarray data. The second method is to make use of on-line resources such as those provided by NLM and NCBI. The functions described in this vignette provide infrastructure for that second type of meta-data usage.

We first describe the functions that allow users to specify queries and open the appropriate web page on their local machine. Then, we investigate the much richer set of tools that are provided by NLM for accessing and working with PubMed data.

2 Using the Browser

There are currently four functions that provide functionality for accessing online data resources. They are:

genbank Users specify GenBank identifiers and can request them to be rendered in the browser or returned in XML.

pubmed Users specify PubMed identifiers and can request them to be rendered in the browser or returned in XML. More details on parsing and manipulating the XML are given below.

locuslinkByID Users specify LocusLink identifiers and the appropriate links are opened in the browser. LocusLink does not provide XML so there is no download option, currently. The user can request that the URL be rendered or returned.

locuslinkQuery Users specify a string that will be used as the LocusLink query and the species of interest (there can be several). The user can request either that the URL be rendered or returned.

Both `genbank` and `pubmed` can return XML versions of the data. These returned values can subsequently be processed using functionality provided by the *XML* package (Temple Lang, 2000). Specific details and examples for PubMed are given in Section 3.

The function `locuslinkByID` takes a set of known LocusLink identifiers and constructs a URL that will have these rendered. The user can either save the URL (perhaps to send to someone else or to embed in an HTML page, see the vignette on creating HTML output for more details).

The function `locuslinkQuery` takes a character string to be used for querying PubMed. For example, this function call,

```
locuslinkQuery("leukemia", "Hs")
```

will find all Human genes that have the word leukemia associated with them in their LocusLink records. Note that the R code is merely an interface to the services provided by NLM and NCBI and users are referred to those sites for complete descriptions of the algorithms they use for searching etc.

3 Accessing PubMed information

In this section we demonstrate how to query PubMed and how to operate on the data that are returned. As noted above, these queries generate XML which must then be parsed to provide the specific data items of interest. Our example is based on the `eset` data from the package *Biobase*. Users should be able to easily replace these data with their own.

```
> library("annotate")
> data(sample.exprSet.1)
> affys <- geneNames(sample.exprSet.1)[490:500]
> affys

[1] "31729_at" "31730_at" "31731_at" "31732_at" "31733_at" "31734_at"
[7] "31735_at" "31736_at" "31737_at" "31738_at" "31739_at"
```

Here we have selected an arbitrary set of 10 genes to be interested in from our sample data. However, `sample.exprSet.1` provided us with Affymetrix identifiers, and for the `pubmed` function, we need to use PubMed ID values. To obtain these, we can use the annotation tools within *annotate*.

```

> library("hgu95av2")
> ids <- getPMID(affys, "hgu95av2")
> ids <- unlist(ids, use.names = FALSE)
> ids <- unique(ids[!is.na(as.numeric(ids))])
> ids

  [1] "7729427" "7835343" "7836461" "7933101" "8121496"
  [6] "8680883" "8764009" "8764062" "8849451" "9054383"
 [11] "9695952" "12878157" "1889752" "8175896" "8325638"
 [16] "9119399" "9765591" "9809674" "10529182" "12408966"
 [21] "12477932" "9199346" "9315667" "9858531" "9858600"
 [26] "11583618" "12679040" "14576168" "14638676" "15563468"
 [31] "15592428" "15710417" "16061479" "1572287" "2005217"
 [36] "2040595" "6548702" "6548703" "7829601" "7916973"
 [41] "7958621" "8735594" "9730618" "10601981" "10750025"
 [46] "11313185" "12198562" "12506116" "12590922" "12614168"
 [51] "14525934" "14681466" "15070778" "15956696" "15956711"
 [56] "15956719" "15956727" "15956728" "15956746" "15967869"
 [61] "16049981" "16381900" "16580233" "7566110" "7566120"
 [66] "9016352" "9221902" "11069162" "12270951" "12885270"
 [71] "14520770" "15331767" "15475563" "1358459" "1973146"
 [76] "2574852" "2878432" "8646877" "9582375" "11857506"
 [81] "7949733" "11076525" "11883959" "11948623" "648524"
 [86] "1074709" "1383230" "1544421" "1808202" "2161834"
 [91] "2395459" "2455687" "2907503" "2936753" "3793760"
 [96] "3842206" "6190133" "6865942" "6893424" "7503723"
[101] "7680654" "7703225" "7780165" "7816144" "7820861"
[106] "7908020" "7908614" "8258301" "8413665" "8422497"
[111] "8567732" "8612804" "8999969" "9199170" "9214391"
[116] "9251807" "9298982" "9417078" "9490638" "9512503"
[121] "9645951" "9668039" "9724539" "9792688" "9832551"
[126] "9852149" "10072583" "10224093" "10366597" "10438535"
[131] "10508519" "10559237" "10564664" "10575001" "10581178"
[136] "10583954" "10601340" "10628374" "10666339" "10793131"
[141] "10882740" "10984435" "11018051" "11027608" "11053353"
[146] "11071908" "11090452" "11106625" "11171322" "11259305"
[151] "11309382" "11333380" "11415434" "11429692" "11438522"
[156] "11461920" "11468283" "11481347" "11524436" "11536049"
[161] "11580270" "11583571" "11701614" "11707283" "11795944"
[166] "11812134" "11821948" "11839310" "11850458" "11877456"
[171] "11889126" "11950878" "11969288" "12006649" "12019270"
[176] "12033495" "12044158" "12048248" "12119014" "12127568"
[181] "12163017" "12383268" "12388543" "12417594" "12466270"
[186] "12475945" "12480535" "12534372" "12571245" "12665555"
[191] "12714565" "12721663" "12849983" "12857853" "14506234"
[196] "14523566" "14576350" "14640694" "14668487" "14684598"

```

```

[201] "14702039" "15072110" "15107502" "15152255" "15198992"
[206] "15236405" "15336687" "15454575" "15468086" "15494313"
[211] "15507486" "15520409" "15679046" "16427282" "1602151"
[216] "7590364" "8619474" "9110174" "12204115" "15302935"
[221] "15308636" "15466936" "15952740" "16130169" "16236267"
[226] "16396496" "10997877" "11448995" "11717165" "11814696"
[231] "12527750" "12631729" "15013951" "15146197" "15185077"

```

We use `getPMID` to obtain the PubMed identifiers that are related to our probes of interest. Then we process these to leave out any that have no PMIDs and we remove duplicates as well. The mapping to PMIDs are actually based on LocusLink identifiers and since the mapping from Affymetrix IDs to LocusLink is many to one there is some chance of duplication. From our initial 11 Affymetrix identifiers we see that there are 235 unique PubMed identifiers (ie papers).

For each of these papers we can obtain information, such as the title, the authors, the abstract, the LocusLink identifiers for genes that are referred to in the paper and many other pieces of information. Again, for a complete listing and description the reader is referred to the NLM website.

We next generate the query and store the results in a variable named `x`. This object is of class `XMLDocument` and to manipulate it we will use functions provided by the XML package.

```

> x <- pubmed(ids)
> a <- xmlRoot(x)
> numAbst <- length(xmlChildren(a))
> numAbst

```

```
[1] 235
```

Our search of the 235 PubMed IDs (from the 11 Affymetrix IDs) has resulted in 235 abstracts from PubMed (stored in R using XML format). The `annotate` package also provides a `pubMedAbst` class, which will take the raw XML format from a call to `pubmed` and extract the interesting sections for easy review.

```

> arts <- vector("list", length = numAbst)
> absts <- rep(NA, numAbst)
> for (i in 1:numAbst) {
+   arts[[i]] <- buildPubMedAbst(a[[i]])
+   absts[i] <- abstText(arts[[i]])
+ }
> arts[[7]]

```

An object of class 'pubMedAbst':

```

Title: Mutations in the carboxy-terminal domain of TBP affect
       the synthesis of human immunodeficiency virus type 1
       full-length and short transcripts similarly.
PMID: 8764009

```

Authors: PS Pendergrast, D Morrison, WP Tansey, N Hernandez
Journal: J Virol
Date: Aug 1996

In the S language we say that the `pubMedAbst` class has a number of different slots. They are:

authors The vector of authors.

pmid The PubMed record number.

abstText The actual abstract (in text).

articleTitle The title of the article.

journal The journal it is published in.

pubDate The publication date.

These can all be individually extracted utilizing the provided methods, such as `abstText` in the above example. As you can see, the `pubMedAbst` class provides several key pieces of information: authors, abstract text, article title, journal, and the publication date of the journal.

Once the abstracts have been assembled you can search them using any of the standard search techniques. Suppose for example we wanted to know which abstracts have the term `cDNA` in them, then the following code chunk shows how to identify these abstracts.

```
> found <- grep("cDNA", absts)
> goodAbsts <- arts[found]
> length(goodAbsts)
```

```
[1] 38
```

So 38 of the articles relating to our genes of interest mention the term `cDNA` in their abstracts.

Lastly, as a demonstration for how one can use the `query` toolset to cross reference several databases, we can use the same set of PubMed IDs with another function. In this example, the `genbank` function is used with the argument `type="uid"`. By default, the `genbank` function assumes that the id values passed in are Genbank accession numbers, but we currently have PubMed ID values that we want to use. The `type="uid"` argument specifies that we are using PubMed IDs (aka NCBI UID numbers).

```
> y <- genbank(ids[1:10], type = "uid")
> b <- xmlRoot(x)
```

At this point the object `b` can be manipulated in a manner similar to `a` from the PubMed example.

Also, note that both `pubmed` and `genbank` have an option to display the data directly in the browser instead of XML, by specifying `disp="browser"` in the parameter listing.

4 Generating HTML output for your abstracts

Many users find it useful to have a web page created with links for all of their abstracts, leading to the actual PubMed page online. These pages can then be distributed to other people who have an interest in the abstracts that you have found. There are two formats for this, the first provides for a simple HTML page which has a link for every abstract, and the other provides for a framed HTML page with the links on the left and the resulting PubMed page in the main frame. For these examples, we will be using temporary files:

```
> fname <- tempfile()
> pmAbst2HTML(goodAbsts, filename = fname)
> fnameBase <- tempfile()
> pmAbst2HTML(goodAbsts, filename = fnameBase, frames = TRUE)
```

5 Session Information

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

```
R version 2.4.0 (2006-10-03)
x86_64-unknown-linux-gnu
```

```
locale:
```

```
LC_CTYPE=en_US;LC_NUMERIC=C;LC_TIME=en_US;LC_COLLATE=en_US;LC_MONETARY=en_US;LC_MESSAGES=en_US
```

```
attached base packages:
```

```
[1] "tools"      "methods"    "stats"      "graphics"   "grDevices"
[6] "utils"      "datasets"   "base"
```

```
other attached packages:
```

XML	GO	Rgraphviz	geneplotter	graph
"1.2-0"	"1.14.1"	"1.12.1"	"1.12.0"	"1.12.0"
xtable	hgu95av2	annotate	Biobase	
"1.4-2"	"1.14.0"	"1.12.1"	"1.12.2"	

References

R. Gentleman and J. Gentry. Querying pubmed. *R News*, 2(2):28-31, June 2002. URL <http://CRAN.R-project.org/doc/Rnews/>.

Duncan Temple Lang. Tools for parsing and generating xml within r and s-plus. *CRAN*, 2000. URL <http://www.omegahat.org/RXML>.

Article Title	Publication Date
Cloning and characterization of hTAFII18, hTAFII20 and hTAFII28: three subunits of the human transcription factor TFIID.	Apr 1995
Generation and initial analysis of more than 15,000 full-length human and mouse cDNA sequences.	Dec 2002
[Expression, purification and characterization of human prorelaxin-like-protein H2 in Escherichia coli]	Sep 2002
Isolation and analysis of the 3'-untranslated regions of the human relaxin H1 and H2 genes.	Apr 2000
Two human relaxin genes are on chromosome 9.	Oct 1984
Relaxin gene expression in human ovaries and the predicted structure of a human preprorelaxin by analysis of cDNA clones.	Oct 1984
Expression of the human relaxin H1 gene in the decidua, trophoblast, and prostate.	Apr 1991
Molecular characterization and pharmacological properties of the human P2X3 purinoceptor.	Jul 1997
Characterization of recombinant human P2X4 receptor reveals pharmacological differences to the rat homologue.	Jan 1997
The HOXC11 homeodomain protein interacts with the lactase-phlorizin hydrolase promoter and stimulates HNF1 alpha-dependent transcription.	May 1998
The human HOX gene family.	Dec 1989
Identification of the alternative splice products encoded by the human protein phosphatase inhibitor-1 gene.	Mar 2002
Identification and characterization of a novel protein inhibitor of type 1 protein phosphatase.	Nov 2000
Isolation and characterization of full-length cDNA clones for human alpha-, beta-, and gamma-actin mRNAs: skeletal but not cytoplasmic actins have an amino-terminal cysteine that is subsequently removed.	May 1983
Isolation and characterization of cDNA clones for human skeletal muscle alpha actin.	Jun 1983
Large-scale concatenation cDNA sequencing.	Apr 1997
Cloning and sequencing of the cDNA encoding human P5.	Oct 1995

Figure 1: pmAbst2HTML without frames
BioConductor Abstract List

The screenshot displays the BioConductor Abstract List interface with frames. On the left, a table lists abstracts with columns for 'Article Title' and 'Publication Date'. The first abstract is selected, and its full text is displayed in a frame on the right. The interface includes the NCBI PubMed logo, a search bar, and various navigation options like 'Limits', 'Preview/Index', and 'History'. The detailed view of the first abstract shows the title 'Cloning and characterization of hTAFII18, hTAFII20 and hTAFII28: three subunits of the human transcription factor TFIID.', authors 'Mengus G, May M, Jacq X, Staub A, Tora L, Chambon P, Davidson I.', and the full text of the abstract.

Figure 2: pmAbst2HTML with frames