

Metadata Harvesting with R and OAI-PMH

Kurt Hornik

2023-01-31

The Open Archives Initiative (<https://www.openarchives.org/>) develops and promotes interoperability standards that aim to facilitate the efficient dissemination of content. One key project is the Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH, <https://www.openarchives.org/pmh/>) which provides “a low-barrier mechanism for repository interoperability” for archives (institutional repositories) containing digital content (digital libraries). OAI-PMH allows people (service providers, such as the ones registered with the OAI listed on <https://www.openarchives.org/service/listproviders.html>) to harvest metadata (from data providers, such as the ones registered with and validated by the OAI listed on <https://www.openarchives.org/Register/BrowseSites/>). Data Providers administer systems that support the OAI-PMH as a means of exposing metadata. Service Providers use metadata harvested via the OAI-PMH as a basis for building value-added services.

OAI-PMH, currently in version 2.0, defines a mechanism for data providers to expose their metadata. The protocol mandates that individual archives map their metadata to the Dublin Core (DC, <https://dublincore.org/>), a simple and common metadata set for cross-domain information resource description. OAI-PMH is a set of six *verbs* or services that are invoked within HTTP, returning the request results in XML format. The OAI-PMH specification can be found at <https://www.openarchives.org/OAI/openarchivesprotocol.html>. Here, we summarize the basic facts and terminology.

A harvester is a client application that issues OAI-PMH requests. A harvester is operated by a service provider as a means of collecting metadata from repositories. Repositories are network accessible servers that can process the six OAI-PMH requests, and are managed by a data provider to expose metadata to harvesters. OAI-PMH distinguishes between three distinct entities related to the metadata made accessible by the OAI-PMH:

resource the object or “stuff” that metadata is “about”. Its nature is outside the scope of the OAI-PMH.

item a constituent of a repository from which metadata about a resource can be disseminated.

record metadata in a specific metadata format. A record is returned as an XML-encoded byte stream in response to a protocol request to disseminate a specific metadata format from a constituent item.

For each item there is an unambiguous unique identifier which is used in OAI-PMH requests for extracting metadata from the item. Items may contain metadata in multiple formats; Dublin Core is mandatory.

Selective harvesting allows harvesters to limit harvest requests to portions of the metadata available from a repository. The OAI-PMH supports selective harvesting with two types of harvesting criteria that may be combined in an OAI-PMH request: timestamps and membership in sets, an optional construct for grouping items.

The XML encoding of records is organized into the following parts:

header contains the unique identifier, a timestamp (the date of creation, modification or deletion of the record), zero or more setSpec elements indicating the set membership of the item, and an optional status attribute for indicating the withdrawal of availability of the specified metadata format for the item, dependent on the repository support for deletions.

metadata a single manifestation (format) of the metadata from an item.

about an optional and repeatable container to hold data about the metadata part of the record. Contents of the containers must conform to an XML Schema. Common uses of these containers include rights statements and provenance statements.

The OAI-PMH verbs (requests) are as follows.

GetRecord retrieve an individual metadata record from a repository.

Identify retrieve information about a repository.

ListIdentifiers an abbreviated form of **ListRecords** which retrieves only headers rather than records.

ListMetadataFormats retrieve the metadata formats available from a repository, or optionally the formats available for a specific item.

ListRecords harvest records from a repository.

ListSets retrieve the set structure of a repository.

Optional arguments to **ListRecords** and **ListIdentifiers** permit selective harvesting of headers based on set membership and/or datestamp.

Some of these requests return a *list* of discrete entities: **ListRecords** returns a list of records, **ListIdentifiers** returns a list of headers, and **ListSets** returns a list of sets. These lists may be large, and it may be practical to partition them among a series of requests and responses. Repositories may reply with incomplete results and a resumption token, which the harvester can use to issue an additional request (and repeat until completion).

The R package **OAIHarvester** provides functions for performing each of the six OAI-PMH requests, using, respectively, packages **curl** (Ooms (2023)) and **xml2** (Wickham, Hester, and Ooms (2021)) for HTTP and XML processing. List requests will automatically be re-issued until complete results are obtained. The names of these verb functions start with ‘**oaih**’ and follow a “combine words with underscores” scheme (e.g., **oaih_list_records**, corresponding to the OAI-PMH **ListRecords** verb, for harvesting records). The functions return the actual (aggregated) *result* of the repository’s response to the harvester’s request.

In addition to these functions for performing OAI-PMH requests, function **oaih_harvest** is a high-level harvester which allows specifying several metadata formats or sets, and giving datestamps as Date or POSIXt date/time objects. Finally, function **oaih_transform** provides functionality for transforming the XML results to “useful” R data structures for further processing or analysis. The results of the verb requests are transformed by default.

The ideas underlying these transformations are best illustrated for harvesting records. In a list context, the result is a list of records, each containing the header (with identifier and datestamp and arbitrarily many set specs), metadata in a certain format, and arbitrarily many about entries. Conceptually, we can think of identifier, datestamp, setSpec, metadata and about as *variables* “observed” for the items in the repository as cases, suggesting the usual rectangular case by variables data organization. When obtaining a single record, it seems natural to transform to a list with these variables. If the rectangular data structure were a data frame, selecting one row (corresponding to a single record) would not straightforwardly yield the single record list transformation (because subscripting list variables in the data frame would give length one sublists rather than the elements). Thus, in the rectangular cases we instead treat rows and columns symmetrically by arranging data in a “list matrix” (a list with a dim attribute, or equivalently, a matrix of list elements). As matrix subscripting drops dimensions when a single row or column is selected, one gets the expected simple list (without a dim attribute) in these cases. (Equivalently, the transformed **oaih_list_records** result is the same as combining the transformed **oaih_get_record** results by rows **rbind**.)

When harvesting records, identifiers and timestamps naturally transform to character strings, and set specs (a header may contain arbitrarily many of these) to character vectors. On the other hand, metadata can be made available in different formats, with different “variables”. We find it more convenient to use a *constant* set of variables for a single transformation of a certain “kind” of OAI-PMH XML results. Thus, we do not immediately transform the metadata, but instead leave them as lists of XML nodes to be transformed in a second stage (with variables differing according to the metadata format; currently, metadata in the Dublin Core and RFC 1807 (<https://www.rfc-editor.org/rfc/rfc1807>) formats can be transformed).

These principles (using lists of single observations on variables and possibly arranging them in a rectangular way, and transforming to constant sets of variables) applies for all transformations of OAI-PMH XML results. Transformations can be added by assigning functions in the (currently internal) environment `oaih_transform_methods_db`.

As an example consider WU Research, an electronic publication platform for research output provided by WU (Wirtschaftsuniversität Wien), which provides an OAI repository at <https://research.wu.ac.at/ws/oai>.

```
> library("OAIHarvester")
> baseurl <- "https://research.wu.ac.at/ws/oai"
```

We can use `oaih_identify` to retrieve information about the repository.

```
> x <- oaih_identify(baseurl)
> rbind(x, deparse.level = 0)

      repositoryName baseURL                                protocolVersion
[1,] "WU Research"   "https://research.wu.ac.at/ws/oai" "2.0"
      earliestDatestamp deletedRecord granularity
[1,] "2022-04-10T02:01:52Z" "no"          "YYYY-MM-DDThh:mm:ssZ"
      adminEmail      compression description
[1,] "pure-support@wu.ac.at" character,0 list,3
```

Here, `rbind` achieves “pretty-printing”: we can see that the repository provides no compression support, and 3 further description entries of kind

```
> vapply(x$description, xml_name, "")

[1] "rightsManifest" "oai-identifier" "Service"
```

where entry 2 indicates that the repository complies with the OAI format for unique record identifiers:

```
> oaih_transform(x$description[[2L]])

$scheme
[1] "oai"

$repositoryIdentifier
[1] "research.wu.ac.at"

$delimiter
[1] ":"

$sampleIdentifier
[1] "oai:research.wu.ac.at:openaire_cris_orgunits/e484dc17-6cd9-476f-85ea-ae8bd2503e40"
```

We can use `oaih_list_metadata_formats` and `oaih_list_sets` to find out about available metadata formats and sets:

```

> oaih_list_metadata_formats(baseurl)

  metadataPrefix
[1,] "oai_cerif_openaire"
[2,] "mods_swepub"
[3,] "oai_dc"
[4,] "mods"
[5,] "nl_didl"
[6,] "qdc"
[7,] "xmetadiss"
  schema
[1,] "https://www.openaire.eu/schema/cris/1.2/openaire-cerif-profile.xsd"
[2,] "http://www.loc.gov/mods/v3/mods-3-7.xsd"
[3,] "http://www.openarchives.org/OAI/2.0/oai_dc.xsd"
[4,] "http://www.loc.gov/mods/v3/mods-3-8.xsd"
[5,] "http://standards.iso.org/ittf/PubliclyAvailableStandards/MPEG-21_schema_files/did/didl.xsd"
[6,] "http://www.openarchives.org/OAI/2.0/oai_dc.xsd"
[7,] "http://files.dnb.de/standards/xmetadiss/xmetadiss.xsd"
  metadataNamespace
[1,] "https://www.openaire.eu/cerif-profile/1.2/"
[2,] "http://www.loc.gov/mods/v3"
[3,] "http://www.openarchives.org/OAI/2.0/oai_dc/"
[4,] "http://www.loc.gov/mods/v3"
[5,] "urn:mpeg:mpeg21:2002:02-DIDL-NS"
[6,] "http://www.openarchives.org/OAI/2.0/oai_dc/"
[7,] "http://www.d-nb.de/standards/xmetadissplus/"

> sets <- oaih_list_sets(baseurl)
> rbind(head(sets, 3L), tail(sets, 3L))

  setSpec          setName          setDescription
"openaire_cris_publications" "OpenAIRE_CRIS_publications" list,0
"openaire_cris_patents"     "OpenAIRE_CRIS_patents"     list,0
"openaire_cris_persons"     "OpenAIRE_CRIS_persons"     list,0
[264,] "openaire_cris_funding" "OpenAIRE_CRIS_funding"     list,0
[265,] "openaire_cris_projects" "OpenAIRE_CRIS_projects"    list,0
[266,] "datasets:all"         "PURE DataSets"             list,0

```

The available formats include the mandatory Dublin Core format, and there is a fairly refined set hierarchy for selective harvesting. To get all publications from year 2005, we can use

```
> x <- oaih_list_records(baseurl, set = "publications:year2005")
```

This gives a “list matrix” with observations of 5 variables on 1902 items:

```

> dim(x)

[1] 1902    5

> colnames(x)

[1] "identifier" "datestamp" "setSpec"    "metadata"  "about"

```

Transforming the Dublin Core metadata is achieved by calling `oaih_transform` on the metadata column, after first removing empty metadata (from deleted records):

```
> m <- x[, "metadata"]
> m <- oaih_transform(m[lengths(m) > 0L])
> dim(m)
```

```
[1] 1902 15
```

giving observations on the 15 (simple) Dublin Core elements:

```
> colnames(m)

 [1] "title"      "creator"    "subject"    "description" "publisher"
 [6] "contributor" "date"       "type"       "format"      "identifier"
[11] "source"     "language"   "relation"   "coverage"    "rights"
```

The topics of the records are available in the ‘subject’ DC variable, with comment “Typically, the subject will be represented using keywords, key phrases, or classification codes. Recommended best practice is to use a controlled vocabulary.” (see <https://dublincore.org/documents/dcmi-terms/#terms-subject>), but without a more detailed syntactic or semantic specification.

Inspecting the output of `m[, "subject"]`, e.g.,

```
> m[head(which(lengths(m[, "subject"]) > 0), 3L), "subject"]
```

```
[[1]]
[1] "/dk/atira/pure/keywords/oefos-6-digit/5/502/502046; name=502046 Economic policy"
[2] "/dk/atira/pure/keywords/oefos-6-digit/5/502/502047; name=502047 Economic theory"

[[2]]
[1] "/dk/atira/pure/keywords/oefos-6-digit/1/102/102000; name=102"
[2] "/dk/atira/pure/keywords/oefos-6-digit/1/102/102001; name=102001 Artificial intelligence"
[3] "/dk/atira/pure/keywords/oefos-6-digit/1/102/102015; name=102015 Information systems"
[4] "/dk/atira/pure/keywords/oefos-6-digit/1/102/102022; name=102022 Software development"

[[3]]
[1] "/dk/atira/pure/keywords/oefos-6-digit/1/102/102000; name=102"
[2] "/dk/atira/pure/keywords/oefos-6-digit/1/102/102001; name=102001 Artificial intelligence"
[3] "/dk/atira/pure/keywords/oefos-6-digit/1/102/102015; name=102015 Information systems"
[4] "/dk/atira/pure/keywords/oefos-6-digit/1/102/102022; name=102022 Software development"
```

shows that “keywords” follow a common scheme of `/dk/atira/pure/keywords` URNs, so we can obtain all keywords via

```
> keywords <- unlist(m[, "subject"])
> keywords <- keywords[!startsWith(keywords, "/dk/atira/pure")]
```

giving a total of 0 keywords. Many of these only occur once:

```
> counts <- table(keywords)
> table(counts)
```

```
< table of extent 0 >
```

The most frequently used keywords are

```
> sort(counts[counts >= 10L], decreasing = TRUE)
```

```
integer(0)
```

showing quite a busy year for the old Research Report Series of the Statistics and Mathematics unit.

To find the records co-authored by myself, one can use

```
> pos <- which(vapply(m[, "creator"],
+                 function(e) any(startsWith(e, "Hornik")),
+                 NA))
```

(note that each creator entry is a character vector of author names). This finds 21 records:

```
> unlist(m[pos, "title"])

[1] "Implications of Probabilistic Data Modeling for Rule Mining"
[2] "Monitoring structural change in dynamic econometric models"
[3] "The design and analysis of benchmark experiments"
[4] "Cluster ensembles"
[5] "No association of clock gene T3111C polymorphism and affective disorders"
[6] "Model-based recursive partitioning"
[7] "Ensemble methods for cluster analysis"
[8] "Open and extensible software for data analysis in management science"
[9] "Additional support for linkage of bipolar disorder to chromosome 3q29"
[10] "A CLUE for CLUster Ensembles"
[11] "Support Vector Machines in R"
[12] "arules - A computational environment for mining association rules and frequent item sets"
[13] "A CLUE for CLUster Ensembles"
[14] "Validation of Credit Rating Systems Using Multi-Rater Information"
[15] "A Lego System for Conditional Inference"
[16] "Residual-based shadings for visualizing (conditional) independence"
[17] "A computational environment for mining association rules and frequent item sets"
[18] "The Strucplot Framework: Visualizing Multi-way Contingency Tables with vcd"
[19] "A Lego System for Conditional Inference"
[20] "Support Vector Machines in R"
[21] "A CLUE for CLUster ensembles"
```

of various types:

```
> table(unlist(m[pos, "type"]))

          article          bookPart contributionToPeriodical
          6          2          1
workingPaper
          12
```

some of which have keywords:

```
> pos <- pos[lengths(m[pos, "subject"]) > 0L]
```

Only one of these is somewhat useful:

```
> unique(m[pos, "subject"])

[[1]]
[1] "/dk/atira/pure/keywords/working_paper_series/1/S12; name=Research Report Series / Department

[[2]]
[1] "/dk/atira/pure/keywords/oefos-6-digit/1/102/102022; name=102022 Software development"
[2] "/dk/atira/pure/keywords/oefos-6-digit/1/101/101029; name=101029 Mathematical statistics"
[3] "/dk/atira/pure/keywords/oefos-6-digit/1/101/101018; name=101018 Statistics"
```

Note that OAI-PMH objects obtained by OAI-PMH requests and subsequent transformations are made up of both character vectors and XML nodes from package **xml2**, with the latter lists of external pointers. Thus, some extra effort is necessary to save OAI-PMH objects to a file or to restore these from a file: see `?oaih_save_RDS` for more information.

References

- J. Ooms. *curl: A Modern and Flexible Web Client for R*, 2023. URL <https://CRAN.R-project.org/package=curl>. R package version 5.0.0.
- H. Wickham, J. Hester, and J. Ooms. *xml2: Parse XML*, 2021. URL <https://CRAN.R-project.org/package=xml2>. R package version 1.3.3.