

Migrating Bibliographic Datasets to the Semantic Web: the AGRIS case

Editor(s): Jens Lehmann, University of Leipzig, Germany; Oscar Corcho, Universidad Politécnica de Madrid, Spain; Richard Cyganiak, DERI Galway, Ireland

Solicited review(s): Name Surname, University, Country

Open review(s): Name Surname, University, Country

Stefano Anibaldi^a, Yves Jaques^a, Fabrizio Celli^a, Armando Stellato^{ab} and Johannes Keizer^a

^a *Food and Agriculture Organization of the United Nations (FAO of the UN), v.le Terme di Caracalla 1, 00154, Rome, Italy*

^b *Department of Enterprise Engineering, University of Tor Vergata, Via del Politecnico 1, 00133, Italy*

Abstract. AGRIS is among the most comprehensive online collections of agricultural and related sciences information. It is a growing global catalog of 4.3 million high-quality structured bibliographic records indexed from a worldwide group of providers. AGRIS relies heavily on the AGROVOC thesaurus for its indexing. Following the conversion of that thesaurus into a SKOS concept-scheme and its publication as Linked Open Data (LOD), the entire set of AGRIS records was also tripled and released as LOD. As part of this exercise, OpenAGRIS, a semantic mashup application, was developed to dynamically combine AGRIS data with external data sources, using a mixture of SPARQL queries and traditional web services. The re-engineering of AGRIS for the semantic web raised numerous issues regarding the relative lack of administrative metadata required to compellingly address the proof and trust layers of the semantic web stack, both within the AGRIS repository and in the external data pulled into OpenAGRIS. The AGRIS team began a process of disambiguation and enrichment to continue moving toward an entity-based view of its resources, beginning with the tens of thousands of journals attached to its records. The evolution of the system, the issues raised during the triplification process and the steps necessary for publishing the result as LOD content are hereby discussed and evaluated.

Keywords: Linked Datasets, Agriculture, Journal Archives, Digital Libraries

1. Introduction

The International Information System for the Agricultural Sciences and Technology (AGRIS) is an initiative set up by the Food and Agriculture Organization of the United Nations (FAO of the UN) in 1974, to make information on agricultural research globally available. One of the main objectives of the AGRIS initiative is to improve access and exchange of information serving “the information needs of developed and developing countries on a partnership basis” [1]. Since 1975, AGRIS has been collecting and disseminating bibliographic information on scientific and socio-economic publications, scholarly papers and grey literature, issued on a wide variety of food and agricultural subjects from over 150 data

providers located in more than 100 different countries.

With the growth of open access institutional repositories (IR), AGRIS has considerably improved its methods for harvesting and indexing metadata from content providers. Initiatives such as the Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) have greatly facilitated interoperability with our data providers.

Today AGRIS disseminates a collection of more than 4.3 million bibliographic references modeled using a qualified Dublin Core metadata format that again ease sharing of information across information systems.

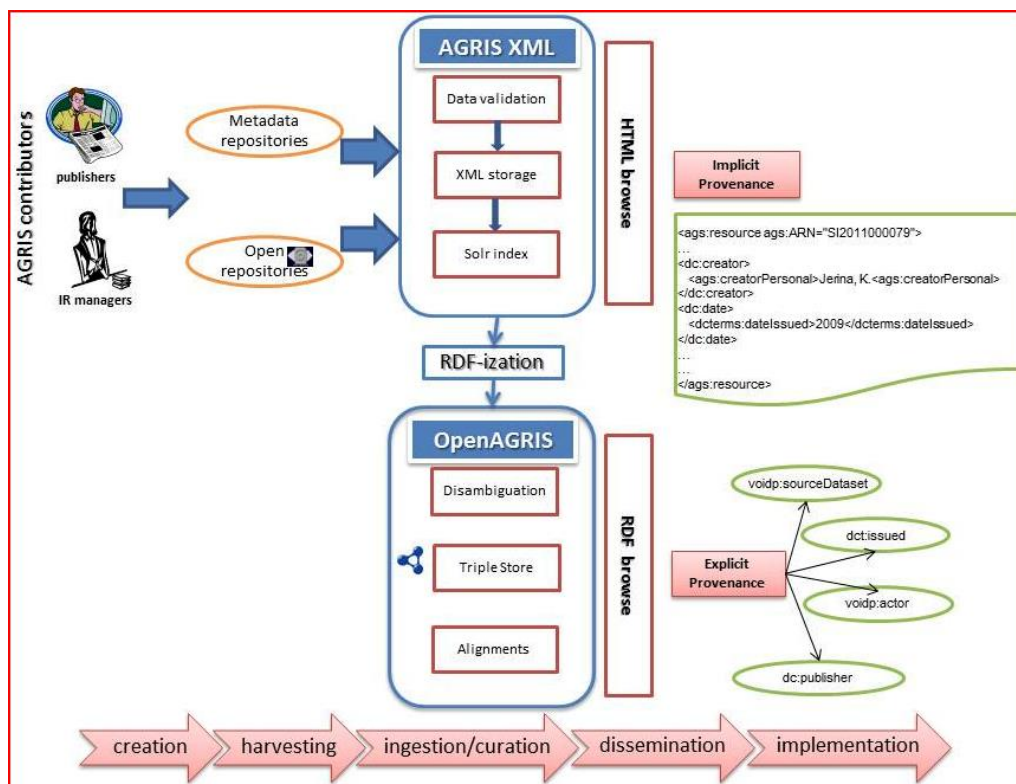


Fig. 1. Derivation history of an AGRIS artifact

Dissemination is handled by the AGRIS portal¹, one of the most trafficked FAO web sites, with over 200,000 unique visitors per month, augmented by the OpenAGRIS [2] semantic mashup that combines and merges information from different online data sources.

The AGRIS data is heavily indexed with the AGROVOC thesaurus. In order to take better advantage of its move to LOD, the entire set of AGRIS records was remodeled using Resource Description Framework (RDF) and deployed in an Allegrograph² triple store. Well-known standard RDF vocabularies such as BIBO [3] FOAF [4] and RDF Dublin Core [5] were reused to facilitate interlinking with other existing datasets. However, the triplication of the AGRIS content brought to light deficiencies in data production, in particular the handling of proof and trust in a world of machine-readable linked data. The paper covers a number of initial issues that were resolved and finishes with an overview of proposals aimed at partially or wholly remedying the remaining proof and trust deficiencies in the AGRIS dataset.

2. Why RDF

Bibliographic resources are static and often do not contain the information needed to answer a user's query, in particular when the record does not link to the full text of the described publication. A recent AGRIS site analysis shows that when end users reach an AGRIS result page (a reference) where no full text information is available, they typically exit the site and presumably search again – using other search engines or databases – for other online resources. The "...digital environment has increased the range of user needs and expectations beyond the scope of the collections of most consortia. End-user services are now required to operate at web-scale and incorporate metadata from sources well beyond the traditional bibliographic record." [6]. The move to RDF has enabled AGRIS to better meet its customers' expectation by providing disambiguated, entity-based access to bibliographic and citation data and by mashing up this information with related data sources. "The ambition is not to collect comprehensively all bibliographic references in the subject area, but to use the latent knowledge in the AGRIS data to find,

¹ <http://agris.fao.org>

² <http://www.franz.com/agraph/allegrograph/>

```

<bibo:Article rdf:about="http://agris.fao.org/aos/records/XS2010X00001">
  < dct:identifier>XS2010X00001 </dct:identifier>
  < dct:title xml:lang="pt">Características anatômicas ...</dct:title>
  < dct:title xml:lang="en">...</dct:title>
  < dct:creator>
    < foaf:Person> < foaf:name>Mesquita, Alessandro Carlos </foaf:name> </foaf:Person>
  </dct:creator>
  < dct:publisher>
    < foaf:Organization> < foaf:name>Instituto Nacional de ... </foaf:name> </foaf:Organization>
  </dct:publisher>
  < dct:issued>2010 </dct:issued>
  < dct:subject rdf:resource="http://aims.fao.org/aos/agrovoc/c_6200"/>
  < bibo:abstract xml:lang="pt"> <![CDATA[As estruturas envolvidas na produção ...]]> </bibo:abstract>
  < bibo:abstract xml:lang="en"> <![CDATA[The structures involved in latex production ...]]> </bibo:abstract>
  < bibo:uri> <![CDATA[http://www.scielo.br/scielo.php?pid=...]]> </bibo:uri>
  < bibo:language>por</bibo:language>
  < dct:isPartOf rdf:resource="http://aims.fao.org/serials/c_e8d916a8"/>
</bibo:Article>

```

Fig. 2: The RDF/XML serialization of an AGRIS record

link and interpret relevant sources on the internet.” [7].

3. The Road to Linked Data

With the exponential growth of open access repositories and web resources, it became clear that the AGRIS objectives required a thorough revision in the strategies and methods to use for such a large dataset.

Following the innovation began by the porting of its indexing thesaurus AGROVOC to a Simple Knowledge Organization System (SKOS) concept-scheme published as LOD [8], the decision was made to connect the AGRIS content to the LOD cloud as well, and to fully exploit the potentialities of the Semantic Web.

3.1. The big move

When the AGRIS team decided to publish its records as linked data in RDF, it quickly became clear that crucial metadata necessary in addressing issues of proof and trust were missing. It goes without saying that “as the number of repositories and aggregators increase, so too do the number of potential formal or informal metadata sources” [9]. Figure 1 shows the long flow of an AGRIS artifact, from genesis to dissemination. Every phase generates administrative metadata, and for each record, AGRIS has

always registered authors, titles, dates and the cataloguing institute.

Becoming part of the LOD cloud meant translating the repository of over 4.3 million XML bibliographic records to RDF, and publishing it on the Web. The translation process required a design step to define vocabularies and properties needed in order to model the data. Keeping in the spirit of reusability, the team avoided almost completely minting any new properties and to reuse as much as possible available standard vocabularies, such as BIBO, FOAF and Dublin Core.

Currently, apart from links to external linked datasets, AGRIS data comprises information from three different RDF datasets maintained in FAO:

- The *AGRIS Records* dataset: the direct translation of AGRIS XML records to RDF. This dataset consists of more than 110 million triples resulting from the conversion of the above mentioned XML repository. The relevancy of the dataset derives not only from the long process of disambiguation and data cleaning, but also on the multilingual assets obtained by crossing its data records with the multiple editions of records available in the diverse countries which provide it with data. Figure 2 shows an example of an AGRIS record.
- The *AGRIS Serials* dataset: since over 80% of AGRIS records are journal articles, it was thought to generate an appropriate authoritative dataset of agricultural journals, with comprehen-

```

<bibo:Journal rdf:about="http://aims.fao.org/serials/c_249cf355">
  <bibo:ISSN>0028-4793</bibo:ISSN>
  < dct:title>New England journal of medicine</dct:title>
  < dct:alternative>The New England journal of medicine</dct:alternative>
  < ags:publisherPlace rdf:resource="http://aims.fao.org/aos/geopolitical.owl#United_States"/>
  < ags:publisherName>Massachusetts Medical Society.</ags:publisherName>
  < dct:language>eng</dct:language>
  < dct:issued>1928</dct:issued>
  < foaf:homepage>http://www.nejm.org/content/index.asp</foaf:homepage>
  < dct:accrualPeriodicity>weekly</dct:accrualPeriodicity>
  < dct:subject rdf:resource="http://aims.fao.org/aos/agrovoc/c_10394"/>
  < dct:subject rdf:resource="http://dewey.info/class/610/about"/>
  < bibo:coden>NEJMAG</bibo:coden>
</bibo:Journal>

```

Fig. 3: The RDF/XML serialization of an AGRIS journal

sive information about each entry (ISSN, subjects, dates, frequency, publisher). The metadata information from the “AGRIS” scientific journals, together with other important data sources such as the journals of the FAO catalogue, the Directory of Open Access journals (DOAJ), Centre for Agricultural Bioscience International³ (CABI) and the National Agricultural Library Catalog⁴ (AGRICOLA) was expanded by retrieving authoritative information from the ISSN Centre database⁵. Currently the team has disambiguated nearly 20,000 agricultural journals (over 90%) and rendered them unique entities stored in a triple store, for a total of almost 320,000 triples. Figure 3 shows a sample of an AGRIS journal in RDF/XML.

- The *Agrovoc* dataset: Agrovoc is the world’s most comprehensive multilingual agricultural vocabulary. Agrovoc contains more than 32,000 concepts in over 20 languages, covering subject fields in agriculture, forestry and fisheries together with cross-cutting themes such as land use, rural livelihoods and food security. A great part of the AGRIS records are indexed with Agrovoc concepts, and the interlinking of this latter to other thesauri (such as Eurovoc⁶, NAL⁷,

DBPedia⁸, etc..) allows for the use of further schemes as indexing elements for AGRIS documents and provides formal crosswalks into other information systems.

3.2. The RDF-ization process

The AGRIS records are stored and indexed on a native XML database, which is the clearinghouse for records (*viz.* Figure 1) which are updated on a monthly basis. The repository is accessed by many applications that pre-date the RDF porting of its content.

For this reason, the translation to RDF requires continuous automatic updates. Moreover, the process is based on operations of data transformation and enrichment, so it requires well defined and structured steps with a precise flow of information (again *viz.* Figure 1).

To begin with, a unique URI is generated for each AGRIS resource: each of the published URIs represents a unique means to identify a specific resource and provides researchers with a partial map of the global research community, linking formal outputs (papers, conference proceedings, etc.) with other information available in other datasets.

³ <http://www.cabi.org/>

⁴ <http://agricola.nal.usda.gov/>

⁵ <http://www.issn.org/>

⁶ <http://eurovoc.europa.eu/>

⁷ <http://agclass.nal.usda.gov/>

⁸ <http://dbpedia.org>

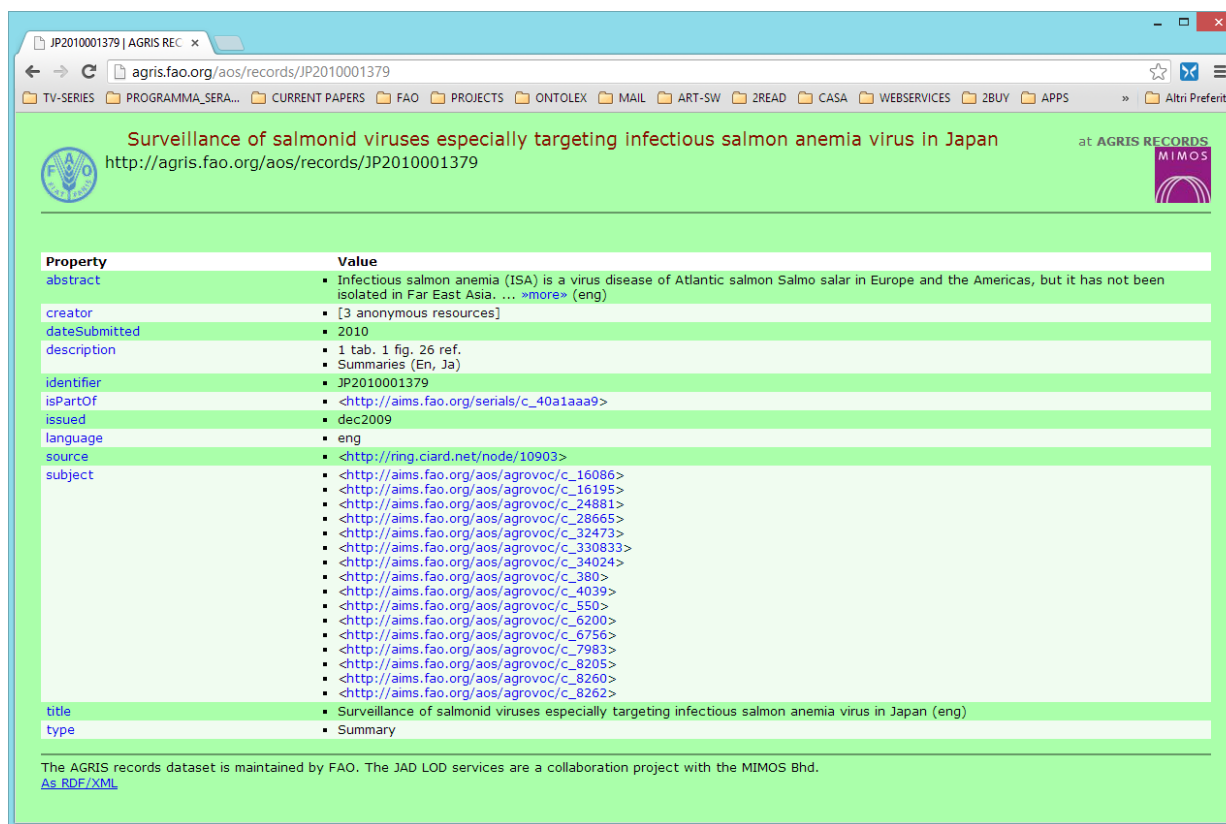


Fig. 4: Description of record JP2010001379, visualized as an HTML page through Pubby

The URI is generated deterministically from available identifiers (see Figure 2, where the identifier associated to *dict:identifier* has been used to also generate the local name of the URI), thus there is no need of a component for managing the consistency of URIs along different translations. Then various steps cover the process of translation of the XML documents to RDF. Figure 1 shows these conversion steps. The process takes as input a set of AGRIS XML records and loads them into the database in an object-relational data model. This step is part of a data validation process necessary to ensure that no duplicate records will be added to the repository and to modify data using the common query language SQL. These data are consumed by a filter that transforms them (data type conversion, data cleaning, data validation) and enriches them with information taken from different controlled and authoritative sources, like the ISSN Centre database and the FAO Open Archive⁹.

⁹ <http://www4.fao.org/faobib/>

At the end, data are ready to be converted to RDF, loaded into a triple store and disseminated to the Web.

This process of RDF-ization is completely automatic: data come directly from structured AGRIS records, so mapping can be between the AGRIS XML and the AGRIS RDF, with a few filters to check the correctness of some information (e.g. date formats, ISSNs, etc...) and to add more data by using predefined data sources that have been previously mapped to AGRIS records. Thus, when new records are added to the AGRIS repository, they are automatically translated to RDF without any loss of reliability as no new uncertainty is added to the dataset. Where critical operations, such as mappings from AGROVOC to other datasets and interactions with data owners, such as the ISSN Centre in Paris, require human intervention, the rest of the translation process is controlled by filters that check data consistency and correctness in the metadata ingested by AGRIS data providers.

Maintaining a wealth of curated knowledge such as the OpenAGRIS Semantic application, entails the implementation of a mechanism for keeping track of provenance information of the data sources. The AGRIS data providers are historically trusted information sources, therefore it was thought to provide provenance information with the objective of helping end users to evaluate the trustworthiness of query results. In addition, proper credits to AGRIS partners are given, specifically when third parties re-publish their information.

3.3. AGRIS Dataset Links

Information about AGRIS is available from the FAO website at:

<http://agris.fao.org/knowledge-and-information-sharing-through-agris-network>

More specific information about the AGRIS Dataset is available at:

<http://aims.fao.org/aos/agrovoc/>

AGRIS data is freely usable under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License.¹⁰

Content negotiation for the LOD Dataset is managed by the server such that clients requesting HTML content (e.g. ordinary web browsers) are given an HTML representation of RDF data describing the requested concept, provided by a Pubby¹¹ installation (see Figure 4), whereas RDF requests result in the return of TDF/XML or TTL.

A description file following the VoID (Vocabulary of Interlinked Datasets) specification [10] is available alongside the AGRIS Linked Open Dataset:

<http://agris.fao.org/void.ttl>

The VoID file contains statistical information about the linked dataset (partitioned into two subsets: AGRIS records and AGRIS journals), as well as coordinates (such as the http address of its SPARQL endpoint) for automatically accessing and properly querying it. More information on the VoID data about AGRIS is reported in section 4.

4. Improving AGRIS Proof and Trust

The issues discussed in the previous section refer to the world of Linked Open Data (LOD) and are typically partially or wholly ignored in closed world systems. In open world systems, data is designed to be accessed by potentially any person or machine, and it is then that problems of proof and trust move to the forefront. OpenAGRIS, as a system designed to access such data, has to cope with these issues and has been an ideal vehicle to get to the top proof/trust layers of the LOD stack to experience the issues firsthand.

Provenance in particular is a broad term that may refer to various levels of granularity. Data provenance was not a historical concern as AGRIS always gave for granted that the agreement with national governments (and the national libraries) provided AGRIS with a license allowing its secretariat to ingest and disseminate data without specific rights or provenance statements. With the shift to digital publishing and machine-readable records however, tracing the provenance chain gained new importance. Dublin Core defines some properties to describe provenance but they are not sufficient in and of themselves to cover all provenance levels. The W3C has recently released two candidate recommendations regarding provenance, offering a data model [11] and ontology [12]. The AGRIS team is following these recent introductions and will soon react to fit its information according to the established vocabularies.

In AGRIS, each record has an identifier known as ARN (AGRIS Record Number), which has a predefined structure and contains information on the data source together with the year of creation of the artifact, namely the bibliographic reference itself. With this unique identifier, AGRIS provides precise and updated statements of the origins of the artifact itself. For instance, the ARN “ES2011001090” refers to a record created in 2011 from a specific AGRIS data provider in Spain, whose progressive number is 1090. Especially for legacy data, ARNs are very important pieces of meta-information, since it is very difficult to retrieve provenance information for decades-old records with poor metadata [13]. Thus, the team tripled information about the AGRIS data providers, providing unique URIs for each and adding triples (using the property *dc:source*) to identify this aspect of provenance, which previously was only implicit in the ARN. More ambitiously, the team plans to move on to the institutions and even the au-

¹⁰ <http://creativecommons.org/licenses/by-nc-sa/3.0/>

¹¹ <http://www4.wiwiw.fu-berlin.de/pubby/>

thors of each record, continuing the process of disambiguation and entity creation.

In fact, provenance must be considered for each piece of metadata and for each triple, since AGRIS knowledge has increased during the RDF conversion of the database, e.g. data such as the unambiguous link to the journal of a publication has been added. Thus, some triples may have a different provenance than the record, and accordingly the previous solution is no longer sufficient as it causes a loss of granularity.

Nevertheless, questions remain on the appropriate set of properties with which to encode such metadata. The Vocabulary of Interlinked Datasets (VoID) has been a likely candidate. Intended as a bridge between publishers and consumers of RDF data, it is organized around four metadata areas: descriptive metadata, access metadata, structural metadata and *linking metadata* that is “helpful for understanding how multiple datasets are related and can be used together” [10]. For descriptive metadata, VoID recommends Dublin Core (DC) and Friend of a Friend properties which together cover basic provenance issues. However to more fully cover provenance, extensions such as the EnAKTing Group’s voidp Vocabulary for Data and Dataset Provenance [14] are desirable while earlier initiatives such as the Open Provenance Model [15] suffer from complex serializations and no reuse of existing properties.

Looking at metadata domains other than provenance, VoID also covers license issues by using DC extensions and contains some recommendations regarding common license types. Where VoID and its extensions are silent is in the trust layer area of quality which will become an important issue as more competing Linked Data comes online and machines need to make dynamic value judgments on which sources to prefer. Though beyond the scope of the proposed work in AGRIS, an interesting initiative which bears future examination is Olaf Hartig’s tRDF, which “proposes a set of criteria to assess the quality of Linked Data sources” [16].

5. The OpenAGRIS Semantic Mashup

In the last two years the AGRIS project has focused its attention on the metadata which end users are mostly interested in and on the many ways to enrich these metadata. AGRIS references often suffer from a lack of complete information and in particular of full text links. Only 4% of the entire collection has

a working full text link. Accordingly, if a user wants to get more information on a specific topic they must use Google or other search engines to retrieve the publication.

The AGRIS team thought then to treat AGRIS records abstractly as metadata sets that could be leveraged to automatically access and display related data. This approach led to the development of OpenAGRIS, a semantic mashup portal that aggregates information from different Web sources using AGRIS records exposed as sets of triples in a Linked Open Data environment. An AGRIS record represented in RDF (see Figures 2, 3) thus becomes the entry point for a mechanism that discovers related web resources primarily via AGROVOC keywords. AGROVOC contains many alignments to other vocabularies (e.g. DBPedia, FAO Geopolitical Ontology, etc.) that allow querying triple stores to retrieve external resources. Moreover, AGROVOC keywords can also be used to query traditional Web Services (e.g. World Bank, FAO fisheries dataset, etc.) to retrieve non-RDF data. The system currently displays production statistics of aquatic species, species occurrence maps, World Bank indicators and more, all dynamically queried through a constellation of related keywords and vocabulary alignments.

OpenAGRIS is an environment that allows the team to test and raise issues related to the implementation of end user systems based on the Semantic Web, and helps in spotting problems in the proof and trust layers of the AGRIS system as well as in most of the systems from which it retrieves data. As previously noted, the first major issue to arise was that of provenance. AGRIS records contained only partial provenance metadata, and most external sources that were accessed contained even less. DBPedia is a typical example in that by its very nature, WIKI data is highly collaborative and almost immune to strong provenance tracking.

The AGRIS triple store currently contains roughly 120 million triples and, ideally, each triple should be linked to administrative metadata. In fact, information about a specific AGRIS record often come from different entities, each of whom have provided part of the metadata during the record lifecycle. It should be possible to know where the information was extracted from, who submitted it and the primary data source. Although this primary source – trustworthy or not – creates the information, it is often only possible to determine the last provider in the chain (e.g. DOAJ).

Licensing is another aspect which is often underestimated and/or kept off the explicit data level. Quality

indicators are also typically non-existent. Finally, there are semantic issues related to vocabulary alignment that can impact the correctness of dynamically retrieved data. A common situation is that two vocabularies representing the same concept with the same name are attached to data that is contextually very different and even wrong from the user's point-of-view. For instance, one vocabulary may have a commercial view of a concept while another has a scientific view. SKOS' deliberately fuzzy definition of *exactMatch* and *closeMatch* properties, while avoiding ontological over-commitment, does little to assist in this regard.

5.1. Details on the OpenAGRIS platform

The core of the OpenAGRIS web application¹² is the AGRIS RDF dataset, stored in an Allegrograph¹³ RDF store, which is the same source exposed on the LOD as a SPARQL endpoint.

By starting from the resource requested by the user and identified by the already mentioned ARN code OpenAgris queries the Agrovoc RDF repository to extract keywords for the specific resource and relationships to other datasets, such as DBpedia. If the resource is a journal article, the engine also queries the AGRIS Serials dataset, obtaining complete information about the journals and related articles from the same journal covering similar topics. This process can be extended to all areas of interest in the linked open data cloud, obtaining all possible information about the specific resource and its main topics.

From an architectural point of view, since filtered SPARQL queries on a triple store are not very efficient, OpenAGRIS uses an Apache Solr index to immediately display information about the record. These indexes, instead of merely providing indexed search over RDF literals, enable semantic shortcuts between elements bound by arbitrary RDF paths, in a way which is specialized for AGRIS constructs. Then, various threads are responsible to query the other linked triple stores by means of their SPARQL endpoints, when available, or other APIs.

6. Conclusions

The process of moving the 4.3 million AGRIS data over to RDF and along the way disambiguating and enriching the records was one of the necessary steps

in the process of rethinking bibliographic databases for the semantic web. Along the way the AGRIS team discovered a number of related benefits, in that the radical openness of the LOD paradigm put into discussion curation practices and shed light on historic weaknesses in the collection of administrative metadata.

The ways in which users search and discover information has changed enormously with the rise of the Internet. It is becoming less and less likely that users are browsing bibliographic data in order to find a print publication that they will then order from a publisher or research library. Users are more and more expecting to simply search subject categories in a generic search engine and click on the results to open full text journal articles, conference proceedings and book excerpts and whatever else will enable them to answer their queries. This implies that traditional bibliographic databases need to find new ways in which to offer value to end users. They must leverage their often rich set of authoritative, disambiguated and interrelated metadata to take advantage of the new sea of external data that is increasingly being exposed via web services and exploited using technologies such as Asynchronous JavaScript and XML (AJAX).

The OpenAGRIS pilot does just that. By combining disambiguated, authoritative records for journals together with a rich set of multi-lingual thesaurus concepts that are in turn broadly linked to other subject thesauri, OpenAGRIS is able to maximize a powerful set of related metadata by using it first against an ever increasing set of semantic web services and then pragmatically against other, useful web services whether semantic or not.

In its original conception OpenAGRIS was to be a purely RDF-based application. Experience showed that only a few domains intersecting with the agricultural domain offered RDF web services. Pragmatic considerations of basic data availability thus necessitated the use of non-RDF APIs.

Despite the relative lack of RDF-based APIs for agricultural use, the field is growing. As new, relevant RDF APIs become available, OpenAGRIS will continue to exploit them, while at the same time doing its part by producing AgriBase, a Freebase-like aggregation API for agricultural data. This challenging endeavor will necessitate wrapping and re-interpreting non-RDF APIs. It will be available via SPARQL endpoint as well as via RDF and non-RDF web services.

Along with a recent rise in the number of RDF-based services, data providers are increasingly

¹² <http://agris.fao.org/openagris/>

¹³ <http://www.franz.com/agraph/allegrograph/>

providing web-service widgets returning complex functionality such as static and dynamic tables, maps, charts and graphs. OpenAGRIS will also begin to exploit such rich services, relying on its unique set of citation-based metadata relationships to deliver relevant information in graphical form to end users.

Finally, behind the scenes, the group responsible for the publication of AGRIS has been quietly publishing and linking an ever increasing set of authoritative RDF vocabularies, from thesauri, to journals, to corporate authors. This growing web of formally linked metadata will continue to improve the relevance of AGRIS records as they reach ever further into the linked data cloud.

References

- [1] Food and Agriculture Organization of the United Nations, "AGRIS introduction," Food and Agriculture Organization of the United Nations, 1981.
- [2] F. Celli, S. Anibaldi, M. Folch, Y. Jaques e J. Keizer, «OpenAGRIS: using bibliographical data for linking into the agricultural knowledge web..» in *Agricultural Ontology Services (AOS)*, Bangkok, Thailand, 2011.
- [3] B. D'Arcus and F. Giasson, "Bibliographic Ontology Specification," 2008. [Online]. Available: <http://bibliontology.com/>.
- [4] "Friend Of A Friend Ontology (FOAF)," [Online]. Available: <http://xmlns.com/foaf/0.1/>.
- [5] DCMI Usage Board, «DCMI Metadata Terms.» 14 June 2012. [Online]. Available: <http://dublincore.org/documents/dcmi-terms/>. [Consultato il giorno 4 March 2013].
- [6] G. Dunsire, «Bibliographic Data Cluster.» 10 February 2011. [Online]. Available: http://www.w3.org/2005/Incubator/ld/wiki/Cluster_BibDat a. [Consultato il giorno 04 March 2013].
- [7] Food and Agriculture Organization of the United Nations (FAO of the UN), "AGRIS 2008-2010: A portal to access global knowledge in agricultural research and technology," Food and Agriculture Organization of the United Nations, 2008.
- [8] C. Caracciolo, A. Stellato, A. Morshed, G. Johannsen, S. Rajbhandari, Y. Jacques e J. Keizer, «Thesaurus Maintenance, Alignment and Publication as Linked Data - The AGROVOC Use Case.» *International Journal of Metadata, Semantics and Ontologies (IJMSO)*, vol. 7, n. 1, pp. 65-75, 14 August 2012.
- [9] T. Emma e J. Allinson, «Signed metadata: method and application.» in *International Conference on Dublin Core and Metadata Applications: "Metadata for Knowledge and Learning"*, Manzanillo, Colima, Mexico, 2006.
- [10] K. Alexander, R. Cyganiak, M. Hausenblas and J. Zhao, "Describing Linked Datasets with the VoID Vocabulary (W3C Interest Group Note)," World Wide Web Consortium, 3 March 2011. [Online]. Available: <http://www.w3.org/TR/void/>. [Accessed 16 May 2012].
- [11] L. Moreau e P. Missier, «PROV-DM: The PROV Data Model (W3C Candidate Recommendation 11 December 2012).» World Wide Web Consortium (W3C), 11 December 2012. [Online]. Available: <http://www.w3.org/TR/prov-dm/>. [Consultato il giorno 6 March 2013].
- [12] T. Lebo, S. Sahoo e D. McGuinness, «PROV-O: The PROV Ontology (W3C Candidate Recommendation 11 December 2012).» World Wide Web Consortium (W3C), 11 December 2012. [Online]. Available: <http://www.w3.org/TR/prov-o/>. [Consultato il giorno 6 March 2013].
- [13] Y. Jaques, S. Anibaldi, F. Celli, I. Subirats, A. Stellato and J. Keizer, "Proof and Trust in the OpenAGRIS Implementation," in *International Conference on Dublin Core and Metadata Applications (DC2012)*, North America, 2012.
- [14] O. Tope e C. Gutteridge, «voidp: A Vocabulary for Data and Dataset Provenance.» EnAKTing, 2011. [Online]. Available: <http://www.enakt.org/provenance/voidp/>. [Consultato il giorno 20 February 2013].
- [15] L. Moreau, B. Clifford, J. Freire, J. Futrelle, Y. Gil, P. Groth, N. Kwasnikowska, S. Miles, P. Missier, J. Myers, B. Plale, Y. Simmhan, E. Stephan e J. Van den Bussche, «The Open Provenance Model core specification (v1.1).» *Future Generation Computer Systems*, July 2010.
- [16] O. Hartig e A. Flemming, «Quality Criteria for Linked Data sources.» 15 December 2010. [Online]. Available: http://sourceforge.net/apps/mediawiki/trdf/index.php?title=Quality_Criteria_for_Linked_Data_sources. [Consultato il giorno 20 February 2013].