This is an electronic version of an article published in Yoose, B., & Perkins, J. (2013). The linked open data landscape in libraries and beyond. *Journal of Library Metadata, 13*(2-3), 197-211 and is available online at:
http://www.tandfonline.com/doi/full/10.1080/19386389.2013.826075

The Linked Open Data Landscape in Libraries and Beyond

Becky Yoose

Grinnell College

Jody Perkins

Miami University

Abstract

For many in the library field, Linked Open Data (LOD) is both a common and an enigmatic phrase. Linked Data has been the topic of many articles, books, conference presentations, and workshops in recent years. The topic, however, is one that many are still working to understand. This article is a brief primer and survey of the current LOD landscape for those who are either new to LOD or wish to reacquaint themselves with LOD. The article will start with a basic introduction to LOD, including some of the standards and formats involved. The second half will describe some of the major LOD projects and efforts in various fields, including libraries, archives, and museums, and provide links to selected resources for those who want to learn more.

Keywords: linked open data, semantic web, metadata, libraries, archives, museums

The Linked Open Data Landscape in Libraries and Beyond

For many in the library field, Linked Open Data (LOD) is both a common and an enigmatic phrase. Linked Data (LD) has been the topic of many articles, books, conference presentations, and workshops in recent years. The topic, however, is one that many are still working to understand, especially in the United States, where adoption of LOD in libraries and other cultural heritage institutions has been slow to develop. As existing library staff take on new roles and new graduates enter the field, they are confronted with the daunting task of learning about Linked Data. Given the limited resources for staff training and development on the job, tackling a topic that is as broad as Linked Data can deter some, and leave others stranded, wondering where they should start.

This article is a brief primer and survey of the current LOD landscape for those who are either new to the subject or wish to reacquaint themselves with LOD. The article will start with a basic introduction to Linked Data and Linked Open Data, including some of the standards and formats involved. The second half will describe some of the major LOD projects and efforts in various fields, including libraries, archives, and museums, as well as provide links to selected resources and tools for those who want to learn more.

LD and LOD – Purpose and Building Blocks

The terms “Linked Data” and “Semantic Web” are sometimes used interchangeably, but while the two terms are connected, they are not one and the same. The Semantic Web is defined as a “Web of Data”. The World Wide Web as we know it today is mostly a “Web of Documents” – individual documents linked to each other – and this Web is geared towards human consumption. The “Web of Data” is what those working with the Semantic Web are striving to

achieve: a Web created for computers linking data that is structured in such a fashion that computers do most of the linking (W3C, 2013c). In order to make the Semantic Web a reality, there needs to be a set of standards and technologies to enable computers to not only read the data in the system but understand it well enough to build relationships between different datasets. This set of standards and technologies is what comprises Linked Data (W3C, 2013a).

The "O" in LOD adds another dimension to LD and the Semantic Web. While the datasets used in LD can have restrictive licensing or in some other way be closed to different uses (Miller as cited in Dulaney, 2012), the Semantic Web cannot reach its fullest potential if the majority of Linked Data have restrictions in place on use and reuse. Tim Berners-Lee created a ranking for LOD sets that gives the first star – the lowest ranking – to any dataset that is openly licensed (for example, a dataset under a Creative Commons license). The rest of the rankings build on that base star, with each step towards full LOD adding another star:

★ Available on the web (whatever format) but with an open license, to be Open Data

★★ Available as machine-readable structured data (e.g. excel instead of image scan of a table)

★★★ as (2) plus non-proprietary format (e.g. CSV instead of excel)

★★★★ All the above plus, Use open standards from W3C (RDF and SPARQL) to identify things, so that people can point at your stuff

★★★★★ All the above, plus: Link your data to other people's data to provide context (Berners-Lee, 2009)

The emphasis on LOD is the openness of the data for reuse, distribution, and/or modification by other systems in the Semantic Web, tying the Open Data movement into the technical standards of LD.

Building Blocks, or What Exactly is Involved in LOD?

The building blocks for LOD are fairly simple. Tim Berners-Lee, in his description of Linked Data, gives four rules in creating LD for the Semantic Web:

1. Use URIs as names for things
2. Use HTTP URIs so that people can look up those names.
3. When someone looks up a URI, provide useful information, using the standards (RDF*, SPARQL)
4. Include links to other URIs. so that they can discover more things. (Berners-Lee, 2009)

Rule one starts with the URI, or uniform resource identifier, which is a string of text that identifies a particular resource on the Internet (W3C, n.d.). This string of text can be broken down into two parts: the protocol needed to access the resource, and the location of the resource. For example, let us look at a mock URI, http://yourlibrary.com/resource/23435. The first part of the URI contains the protocol that should be used to retrieve the resource; in this case, the hypertext transfer protocol (HTTP) is used. This protocol will then bring you to the resource located at the address yourlibrary.com/resource/23435 (assuming that the resource exists, can be accessed, and is not offline). Other common protocols one might see in an URI are the File Transfer Protocol (FTP, used for file transactions between machines) and mailto. There has been confusion about the relationship between URI and URL; however, an URL (uniform resource locator) simply is a form of an URI (URI Planning Interest Group, W3C/IETF, 2001). This is the

main gist of rule two, which states that HTTP should be used in URIs to make them discoverable to other people. Though it might be an unspoken assumption, URIs should be unique, and should not be used multiple times to describe different resources. There are many unique identifiers that the libraries, archives, and museums (LAM) fields incorporate into their metadata. However, there are identifiers that many assume are unique, but in practice are not. While it may be tempting to use the ISBN as part of an URI for that resource, the fact that ISBNs have been reused by publishers makes the ISBN a nonunique identifier, making it inappropriate to use the ISBN as the unique identifier in the URI address section.

Once an HTTP URI (URL) has been identified for a resource, we move on to rule three: using standards to make the information located at the URI useful, not only to the person looking at the information, but also to machines, which need a standardized structure to use the data for various applications and processes. There are many technologies and standards listed in the LD suite, but there are a few key standards that are used throughout the LD landscape. The standard that is referred to the most when many discuss LD is the Resource Description Framework (RDF). RDF is a data model that, through XML syntax, represents resources that reside on the Web (Manola & Miller, 2004). The core of an RDF statement is the triple, which is made up of a subject, a predicate, and an object.

In the English language, basic sentences contain these three things; for example, the sentence “Mary has a pet that is a little lamb” states that Mary (the subject) owns a pet (the predicate), which happens to be a lamb (the object). This sentence can be represented as an RDF triple like this:

<http://example0.com/mary>

<http://example1.com/term/pet><http://example2.com/term/animal/lamb>

Namespaces can be used to make RDF triples more readable:

<example0:mary> <example1:pet> <example2:lamb>

To retrieve RDF data, there are several standards and protocols available for use. The standard that is used most commonly in retrieving RDF data is Simple Protocol And RDF Query Language (SPARQL), a language that can be used to search and retrieve RDF results via HTTP or SOAP (W3C, 2013b). SPARQL is similar to other querying languages, like SQL, but is designed specifically to work with RDF triples, like the one shown in the example above. An example SPARQL query to see who has a pet lamb shows the similarity of SPARQL to SQL syntax:

```
PREFIX m: http://example1.com/term/
PREFIX n: http://example2.com/term/animal
SELECT ?person
WHERE
{ ?person m:pet n:lamb . }
```

The results from a SPARQL query can be delivered in three formats: XML, CSV/TSV, and JSON. This choice in results formatting gives flexibility in application development when manipulating the data from a particular query.

Another standard used in LOD is OWL Web Ontology Language. OWL is a language for use in the Semantic Web to express knowledge of and relationships between things through the use

of ontologies (Hitzler, Krötzsch, Parsia, Patel-Schneider, & Rudolph, 2012). OWL uses the following in its modeling of this knowledge:

- Axioms: the basic statements that an OWL ontology expresses
- Entities: elements used to refer to real-world objects
- Expressions: combinations of entities to form complex descriptions from basic ones (Hitzler, Krötzsch, Parsia, Patel-Schneider, & Rudolph, 2012)

OWL ontologies are hierarchical in structure and consist of instances (or objects), classes (which can contain a group of instances), and properties (which give characteristics of classes). RDF and XML both lack a robust vocabulary to express relationships. However, OWL allows for the expression of relationships between classes and properties, as well as providing a way for more detail in property characteristics. An OWL ontology can be expressed in various serializations, including RDF triples and XML. The RDF/XML example below involves an instance (the lamb) and a class (animal):

```
<Animal rdf:about="Lamb"/>
```

If one wanted to indicate Mary is the owner of the lamb, then the RDF statement would need to include a property (in this case, hasOwner):

```
<rdf:Description rdf:about="Lamb">
    <hasOwner rdf:resource="Mary"/>
</rdf:Description>
```

By using the structures and guidelines provided by the OWL standards documentation, the data created and used in the Semantic Web are represented similarly regardless of which ontologies were created and/or used in the expression of that data. In addition, OWL can be used in

conjunction with other data models, like we see with the RDF example above, which can be helpful when a more robust and formal data model is needed.

One more standard that has been used in the LAM environment is Simple Knowledge Organization System (SKOS), a data model that can be used to express different types of controlled vocabularies (Isaac and Summers, 2009). The structure of SKOS is such that it can be used to model existing controlled vocabularies and ontologies for use in the Semantic Web. A real world example of the use of SKOS in LAM is the Linked Data services provided by the Library of Congress at http://id.loc.gov. These services provide LOD versions of the Library of Congress name and subject authorities, along with various MARC code sets. According to Library of Congress, using SKOS in lieu of creating a new XML schema (using OWL) gave them access to tools that allowed for translating existing data into LD (Technical Center, n.d.).

Why Should Libraries Care, or the Purpose of LOD in LAM

In his article "Library Linked Data Now!" from 2009, Ross Singer argued for the widespread adoption of LD in the library field, noting that the metadata created and maintained by libraries would benefit both libraries themselves and the greater information community on the Web. Since then there has been growth in LOD in the field, though some would argue that it has been slower than desired (Alemu, Stevens, Ross, & Chandler, 2012). Though the adoption of LD/LOD in libraries has been slow to some, the number of pilot projects is growing, a body of literature is emerging and working groups have formed around possible uses for LOD. Before diving into LOD one might ask why LAM institutions should care about LOD. What is the purpose of doing LOD in LAM?

In essence, LOD allows for the structured metadata created and maintained by LAM institutions to be shared in such a way that the general community can interact and enrich the data. Otherwise, it is very difficult to retrieve most of the formats that LAM institutions use to store and access their metadata. In addition, most of the formats and standards that LAM institutions use to store and access their metadata are unfamiliar to non-LAM users. LOD is a framework of standards and technologies that can complement existing practices, vocabularies, and ontologies used in LAM. Libraries can use existing LAM standards while creating and publishing LD (Alemu, Stevens, Ross, & Chandler, 2012). Additionally, there are LD vocabularies specific to libraries, which can help expose library metadata. An example of such a set of vocabularies is the RDA Vocabularies (http://rdvocab.info/). The Vocabularies allow those who have created library metadata under the RDA standard to format the metadata in RDF using elements that are closely tied to RDA. Going back to the "Web of Data" idea mentioned above, LAM metadata is a considerable contribution to the creation and curation of that Web.

While there are large-scale examples of LOD in LAM that will be covered below, there have been libraries that have implemented LD/LOD in their local systems. Several examples can be found in the Code4Lib Journal, including examples in metadata conversion and LD search functionality. Some examples are more experimental, as in the case of Westrum, Rekkavik, and Talleras (2012). The authors took a subset of FRBRized MARC records representing the work of two particular authors and converted them to RDF. From there, the metadata was linked to data in the LOD cloud, making it possible for the library to write applications for library users that could take advantage of this new enriched metadata. In another example of enriching metadata, Holgersen, Preminger, and Massy (2012) converted user-generated content, such as reviews and

tags, to LOD making it available to other libraries for use in their own catalogs. In addition to enriching library metadata, others have delved into serving LOD to constituents, as in the case of Johnson's (2013) example of using JSON-LD, BibJSON, and Elasticsearch for creating a robust index that can incorporate multiple LOD sets.

Others have proposed various ways that LOD can be used in specific parts of LAM. Krier (2012), in her article on using FRBR and RDF for cataloging serials, points out various benefits of using LOD formats over other traditional library metadata formats. An example of a benefit deals with how different formats represent journal families. While the MARC 76X-78X fields have been used to link a journal to a particular family, the benefit of using RDF and LOD would ensure that a link that would otherwise have been misrepresented or not represented fully in a local system would have other LOD sets to show the complete journal family links. There are also other major implications for libraries. Schreur (2012) discusses ways in which LD has changed and will continue to change the ways that libraries perform authority control and determine subject access for library resources, as well as more general paradigm shifts in the way that libraries interact with library metadata. For example, one of the shifts would be from creating and maintaining individual records to linking individual statements, which would have major implications for the way that technical services departments currently operate.

Doing LOD - Next Steps

The LOD landscape is vast, and LOD is versatile as demonstrated above. For many in the library world who want to step into this landscape, this vastness and versatile nature can lead to hesitation and confusion. Where should one start? Some might consider transforming their main bibliographic database of MARC records the first place to start; however, the size of the

database, as well as the complications that come with transforming a large amount of MARC data into suitable LOD, can be a daunting task, leading to frustration and eventual abandonment of the project. As with learning anything new, starting small makes the learning process more manageable. In the case of LOD, Byrne and Goddard suggest starting with a subset of data from a small collection (2010). With this in mind, the following are commonly available tools and resources that one could use in their first steps to create and publish library LOD.

Drupal

Drupal, an open source content management system, is widely used in many libraries. Libraries have used Drupal in various ways, including hosting a library website, integration of library catalogs, and building digital repositories. Starting with version 7, the Drupal core includes support for publishing RDF and RDFa (a set of attributes that allow for the addition of machine-readable data in human readable data) . To make full use of this RDF support in Drupal 7, one might be advised to go beyond the standard RDF module; luckily, the active development community in Drupal has many modules that one can choose from, as well as groups that focus on using Drupal in the Semantic Web (Stevenson, 2011; Semantic Web: Drupal groups, n.d.). If there is an existing standalone collection on Drupal, the RDF support and supporting modules would be a first step to publishing the collection data as LOD.

Open Refine

Open Refine, formerly Freebase Gridworks and Google Refine, is an open source application that has many uses: it can clean messy data, standardize it, link the data to other public databases, and export it in the structure of one’s choosing. If one has, for example, a spreadsheet of metadata that was exported from a legacy database, Open Refine can take that

spreadsheet through most of the steps needed to restructure that metadata into LOD. While exporting data as RDF is not within the standard Open Refine toolset, there is a RDF extension that can be added to Open Refine, which will allow one to not only export data in RDF but also incorporate existing vocabularies or custom build a vocabulary for the RDF dataset (DERI Linked Data Research Centre, n.d.).

While Open Refine is a very useful tool in turning existing datasets into LOD, it does involve more work than installing and configuring the Drupal modules mentioned above. For guidance in cleaning and structuring data as LOD, Free Your Metadata has tutorials and guides to walk one through the process, as well as a listing of upcoming events and workshops (Multimedia Lab & MaSTIC, 2013). In addition, Open Refine has a list of other tutorials and guides that might be relevant to the library field (External Resources, 2013).

Registries and Hubs

The Data Hub

http://datahub.io/

The Data Hub was established in 2007 as an open registry and catalog of datasets (most openly licensed). It uses CKAN, an open source platform developed and maintained by the Open Knowledge Foundation (OKF). The Hub is community-run and groups can be set-up around specific datasets or community interests such as Bibliographic Data, Open Archeology, Climate Data, etc. It includes information for 6,000 data sets in 23 languages. Entries include details on licensing, owner contact, size (in triples), links to namespaces and SPARQL endpoints. The Linked Open Data Cloud is drawn from 295 community maintained entries in the Hub. Approximately one third of these are described using VoID http://www.w3.org/TR/void/ - an

RDF Schema vocabulary which includes descriptive (Dublin Core), access, structural and relational elements.

Open Metadata Registry

http://metadataregistry.org/

Intended to be part of a distributed registry system the OMR is an open registry of metadata schemas (element sets), schemes (controlled vocabularies), and application profiles, developed to support reuse, standardization and interoperability of metadata. Originally established in 2005 to support the work of the National Science Digital Library it took on a much broader scope in 2010 and continues independent of NSDL. It can provide contributors with namespace services and the means to automate creation and maintenance of element sets and application profiles. Contributors can submit schemas and schemes to a registry workflow for review and publication. It currently contains 69 schemas and 296 schemes.

Linked Open Vocabularies

http://lov.okfn.org/dataset/lov/

LOV is an open registry containing entries for 348 vocabularies (RDFS or OWL) used or useable in the Linked Open Data Cloud. Originally part of the Datalift Project, LOV is maintained by Pierre-Yves Vandenbussche and Bernard Vatant. It was adopted as an official project of the Open Knowledge Foundation in 2012. One of the goals as stated on the project website is to make "… explicit the ways they [vocabularies] link to each other and provide metrics on how they are used …" using VOAF, a specification for defining properties that express such dependencies. LOV is searchable at the element or vocabulary level, and includes

information on ownership, licensing, and namespaces, as well as statistics on distribution and use. The LOV dataset is licensed under Creative Commons CC BY 3.0.

DBpedia

http://www.wiki.dbpedia.org

Started in 2007 by folks at the Free University of Berlin, the University of Leipzig and OpenLink Software, DBpedia is currently part of the W3C Linked Open Data community project. It is not a registry but a knowledgebase of structured data extracted from Wikipedia and served as Linked Data with dereferenceable URIs. It covers a broad range of domains and includes identifiers for 3.64 million concepts. The dataset is comprised of 1.89 billion RDF triples: 400 million are from Wikipedia (English edition), 1.46 billion from other language editions, and 27 million link to external RDF datasets. Over a period of almost four months in 2009 DBpedia URIs received 561,277 hits per day and SPARQL endpoints received 177,734 queries per day. The DBpedia Wiki includes a number of use cases and applications written to leverage DBpedia data. For access to the complete dataset see http://wiki.dbpedia.org/OnlineAccess

Library Initiatives

The Library of Congress Bibliographic Framework Initiative

http://www.loc.gov/bibframe/.

The Library of Congress Bibliographic Framework Initiative is laying the groundwork for libraries to make a transition into the Linked Data world. It is being developed as a high-level LOD model that can be used to transform MARC21 records into Linked Data. The

The BIBFRAME model splits MARC21 into 4 main classes – Work, Instance, Authority and Annotation. Chief among its goals is to be content model (RDA, DACS, CCO) agnostic. However, at this stage RDA is a prominent content type. The BIBFRAME vocabulary is being published as a single namespace and the model is defined in RDF. Transforming the MARC21format into a Linked Data Model will be followed by an extended period of testing. Work is underway to develop tools and support services. Early experimenters include the British Library, George Washington University, Princeton University, Deutsche National Bibliothek, National Library of Medicine, OCLC, and the Library of Congress. The BIBFRAME Model as well as experimental tools and services are available at http://bibframe.org/

The Library of Congress Linked Data Service

http://id.loc.gov/

Since 2009 the Library of Congress has been publishing LC name authorities, subject headings and other LC vocabularies and standards as LOD. URIs have been issued for the vocabularies and the individual values they contain. Currently there are 17 datasets available including Library of Congress Subject Headings and LC classification schedules – B (Philosophy, Psychology, Religion), K (Law), M (Music), N (Fine Arts) and Z (Bibliography, Library Science, Information Resources) with the remainder to be released incrementally.

Other National Libraries have been active in the LOD domain as well. **The British Library** will release a version of the British National Bibliography as LOD including both monographic and serial publications http://www.bl.uk/bibliographic/datafree.html. **The German National Library** is partnering with many other institutions on LOD projects and has

established their own Linked Data service including both bibliographic and authority data http://www.dnb.de/EN/Service/DigitaleDienste/LinkedData/linkeddata_node.html

VIAF: Virtual International Authority File

http://viaf.org

VIAF started as an OCLC research project in 2003 jointly with the Library of Congress, the Deutsche Nationalbibliothek, the Bibliotheque Nationale de France and an expanding list of other national libraries. In 2012, it transitioned from research status to an OCLC international service. There are currently 20 contributing agencies from 16 countries. VIAF links authority records for named entities from different regions and countries together to create a merged super authority record cluster, allowing variations of authorized forms to co-exist. There are currently 20 million clusters, each with a URI, available as LOD. OCLC has also released FAST: Faceted Application of Subject Terminology as Linked Data. Both are available by OCLC under the ODC-By license.

Archive and Museum Initiatives

SNAC: Social Networks and Archival Context Project

http://socialarchive.iath.virginia.edu/

Funded by the National Endowment for the Humanities and the Mellon Foundation, SNAC is a collaboration between the Institute for Advanced Technology in the Humanities at the University of Virginia, the University of California at Berkeley School of Information, and the California Digital Library. The project will leverage the content of existing archival descriptions contributed by a wide range of institutions in the United States and Europe. It uses an automated process to create EAC-CPF (Encoded Archival Context-Corporate Bodies, Persons, and

Families) records, part of the EAD family of standards - extracting named entities from collection descriptions, normalizing, enhancing and linking them to related resources. The project aims to pave the way for integrated access to archival resources, revealing connections and shedding light on their socio-historical context. It plans to make EAC-CPF records available as Linked Open Data using RDF mapping in the final phase of the project.

LOCAH: The Linked Open Copac and Archives Hub Project

http://data.archiveshub.ac.uk/

LOCAH is a JISC funded project to make Archives Hub and Copac data available as structured Linked Data. Copac brings together the catalogues of over 70 major UK and Irish libraries. The Archives Hub is a gateway to thousands of the UK's richest archives and represents over 220 institutions. The vision for the project puts "archival and bibliographic data at the heart of the Linked Data Web, enabling new links between diverse content sources and the free and flexible exploration of data so that researchers can make new connections between subjects, people, organizations and places to reveal more about our history and society." Project partners include Eduserv, Talis and OCLC**.** **Linking Lives** is an extension of LOCAH and explores ways to present biographical Linked Data for research purposes, drawing on datasets from external sources to make connections and provide additional context.

Smithsonian American Art Museum

In partnership with the Information Sciences Institute and the Department of Computer Science at the University of Southern California, the Smithsonian's American Art Museum will make its entire database of 40,000 records available as 5-star Linked Open Data. Intended to be a model project for other museums, five phases have been outlined: data preparation, ontology

development, RDF mapping, linking to hub datasets, and publishing the data. The ontology draws from the Europeana Data Model, SKOS, Dublin Core, RDA and schema.org. Mapping was done with KARMA, an open source data integration tool developed by ISI. Links have been made to external hubs such as DBpedia (2,194) and the New York Times (70) with plans to link to the Union List of Artists Names, GeoNames, social media sites and other museums. The final dataset will be published under a CC0 license. Future plans include the development of a tool enabling users to curate stories using LOD. The project does not currently have an official website (see Szekely, et al, 2013 and Goodlander, G., 2013, May 6).

Europeana

http://data.europeana.eu

Europeana, the flagship digital library project in the European Union, collects metadata from approximately 1,500 cultural institutions, providing access to millions of cultural resources including books, films, museum objects and archival records that have been digitized across Europe. Development of the Europeana Data Model in 2010 marked the beginning of a serious commitment to LOD. The data.europeana.eu project started with a limited number of data providers as an experimental pilot in February 2012. Harvesting 2.4 million items from 8 Europeana contributors and 200 institutions across 15 countries, and resulting in approximately 185 million triples - it is one of the largest Linked Open Datasets available in the cultural heritage community. The current version includes metadata for 20 million texts, images, videos and sounds. Datasets are available under a CC0 license.

Conclusion

It is still early days for LOD and there is still much to be done. LOD is moving things to global scale as a number of the initiatives and projects above demonstrate. Openness, collaboration and cooperation will bring progress and carry us forward into the world of Linked Open Data. As archivists interested in this emerging technology will tell you, our history, our stories are all connected. Libraries, archives and museums have accumulated an embarrassment of riches in the form of unique digitized resources and structured data as well as un-mined unstructured content, all of which are lying fallow inside a Web of documents and untapped relationships. Failure to embrace the possibilities of the Semantic Web would be to tragically squander not only those riches but also all the resources expended to create them.

References

Alemu, G., Stevens, B., Ross, P., & Chandler, J. (2012, May 24). Linked data for libraries: benefits of a conceptual shift from library-specific record structures to RDF-based data models. *78th IFLA General Conference and Assembly*. Retrieved from http://conference.ifla.org/past/ifla78/92-alemu-en.pdf

Berners-Lee, T. (2009, June 18). Linked data - design issues. Retrieved from http://www.w3.org/DesignIssues/LinkedData.html

The British Library. (n.d.). Free Data Services. Retrieved from http://www.bl.uk/bibliographic/datafree.html.

Byrne, G., & Goddard, L. (2010). The Strongest Link: Libraries and Linked Data. D-Lib Magazine, 16(11/12), 2. doi:10.1045/november2010-byrne

Datahub. (n.d.). The Datahub. Retrieved from http://datahub.io/

DBpedia. (2013). DBpedia. Retrieved from http://wiki.dbpedia.org

DERI Linked Data Research Centre. (n.d.). GRefine RDF Extension. Retrieved from http://refine.deri.ie/

Dulaney, C. K. (2012). Libraries, Linked Data, and the Semantic Web: Positioning Our Catalogs to Participate in the 21st Century Global Information Marketplace: An ALCTS Midwinter Symposium; ALA Midwinter Meeting, Dallas, TX, January 20, 2012. *Serials Review*, *38*(2), 159-160. doi:10.1016/j.serrev.2012.05.008

Europeana. (2012). Linked Open Data. Retrieved from http://data.europeana.eu

External Resources. (2013, April 13). Retrieved May 16, 2013 from the OpenRefine Wiki: https://github.com/OpenRefine/OpenRefine/wiki/External-Resources

FAST. (2013). OCLC FAST Linked Data. Retrieved from http://fast.oclc.org/fast/

The German National Library. (2013). Linked Data Service of the German National Library. Retrieved from http://www.dnb.de/EN/Service/DigitaleDienste/LinkedData/linkeddata_node.html

Goodlander, G. (2013, May 6). Linked Open Data and American Art. *Presentation for the Linked Open Data American Art Collaborative on April 29, 2013*. Retrieved from http://www.slideshare.net/georginab/linked-open-data-and-american-art

Hitzler, P., Krötzsch, M., Parsia, B., Patel-Schneider, P. F., & Rudolph, S. (eds.). (2012). OWL 2 web ontology language primer (second edition). Retrieved from http://www.w3.org/TR/owl2-primer/

Holgersen, R., Preminger, M., & Massey, D. (2012). Using Semantic Web Technologies to collaboratively Collect and Share User-Generated Content in Order to Enrich the presentation of Bibliographic Records-Development of a Prototype Based on RDF, D2RQ, Jena, SPARQL and WorldCat's FRBRization Web Service. *Code4lib Journal*, (17), 1-17. Retrieved from http://journal.code4lib.org/articles/6695

Isaac, A., & Summers, E. (eds.). (2009). SKOS simple knowledge organization system primer. Retrieved from http://www.w3.org/TR/skos-primer/

Johnson, T. (2013). Indexing Linked Bibliographic Data with JSON-LD, BibJSON and Elasticsearch. *Code4lib Journal*, (19), 1-11. Retrieved from http://journal.code4lib.org/articles/7949

Krier, L. (2012). Serials, FRBR, and Library Linked Data: A Way Forward. *Journal Of Library Metadata*, *12*(2/3), 177-187. doi:10.1080/19386389.2012.699834

Library of Congress Bibliographic Framework Initiative. (n.d.). Bibliographic Framework Initiative. Retrieved from http://www.loc.gov/bibframe/

Library of Congress Bibliographic Framework Initiative. (2013). BIBFRAME.ORG: Bibliographic Framework Initiative – Overview. Retrieved from http://bibframe.org/

Library of Congress Linked Data Service. (2013). LC Linked Data Service. Retrieved from http://id.loc.gov/

Linked Open Vocabularies. (2013). Linked Open Vocabularies (LOV). Retrieved from http://lov.okfn.org/dataset/lov/

LOCAH. (2010). LOCAH: The Linked Open Copac and Archives Hub Project. Retrieved from http://data.archiveshub.ac.uk/

LOCAH. (2012). LOCAH Project Blog. Retrieved from http://archiveshub.ac.uk/locah/

Manola, F., & Miller, E. (eds). (2004). RDF primer. Retrieved from http://www.w3.org/TR/rdf-primer/

Multimedia Lab, & MaSTIC. (2013). Free your metadata. Retrieved from http://freeyourmetadata.org

Open Metadata Registry. (2013). The Registry. Retrieved from http://metadataregistry.org/

OpenRefine. Retrieved from http://openrefine.org/

Schreur, P. (2012). The Academy Unbound. *Library Resources & Technical Services*, *56*(4), 227-237.

Semantic web: Drupal Groups. Retrieved from http://groups.drupal.org/semantic-web

Singer, R. (2009). Linked Library Data Now!. *Journal Of Electronic Resources Librarianship*, *21*(2), 114-126.

SNAC (n.d.). SNAC: the Social Networks and Archival Context Project. Retrieved from http://socialarchive.iath.virginia.edu/

Stevenson, K. (2011, August 31). How does RDF work in Drupal 7? Retrieved from http://www.lullabot.com/articles/how-does-rdf-work-drupal-7

Szekely, P., Knoblock, C., Yang, F., Zhu, X., Fink, E., Allen, R. & Goodlander, G. (2013). Connecting the Smithsonian American Art Museum to the Linked Data Cloud. *In the Proceedings of the10th Extended Semantic Web Conference, Montpellier, May 2013.* Retrieved from http://www.isi.edu/~szekely/contents/papers/2013/eswc-2013-saam.pdf

Technical Center. Retrieved from http://id.loc.gov/techcenter/metadata.html

URI Planning Interest Group, W3C/IETF. (2001, September 21). URIs, URLs, and URNs: clarifications and recommendations 1.0. Retrieved from http://www.w3.org/TR/uri-clarification/

VIAF. (2012). VIAF (the Virtual International Authority File). Retrieved from http://viaf.org

W3C. (2011). Describing Linked Datasets with the VoID Vocabulary. Retrieved from http://www.w3.org/TR/void/

W3C. (2013a). Linked data. Retrieved from http://www.w3.org/standards/semanticweb/data

W3C. (2013b). Query. Retrieved from http://www.w3.org/standards/semanticweb/query

W3C. (2013c). Semantic web. Retrieved from http://www.w3.org/standards/semanticweb/

W3C. (n.d.). Web naming and addressing overview (URIs, URLs,...). Retrieved from http://www.w3.org/Addressing/

Westrum, A., Rekkavik, A., & Talleråas, K. (2012). Improving the presentation of library data using FRBR and Linked data. *Code4lib Journal*, (16), 1-7. Retrieved from http://journal.code4lib.org/articles/6424

Appendix A
LOD Resources

Linked Data
http://linkeddata.org/

Semantic Browsers
http://www.w3.org/wiki/TaskForces/CommunityProjects/LinkingOpenData/SemWebClients

Semantic Search Engines
http://www.w3.org/wiki/TaskForces/CommunityProjects/LinkingOpenData/SemanticWebSearchEngines

Guides and Tutorials
http://linkeddata.org/guides-and-tutorials

W3C Library Linked Data Incubator
http://www.w3.org/2005/Incubator/lld/