# The Landscape of Ontology Reuse in Linked Data

María Poveda-Villalón, Mari Carmen Suárez-Figueroa, Asunción Gómez-Pérez

Ontology Engineering Group. Departamento de Inteligencia Artificial.
Facultad de Informática, Universidad Politécnica de Madrid. Spain

{mpoveda, mcsuarez, asun}@fi.upm.es

**Abstract.** The uptake of Linked Data (LD) has promoted the proliferation of datasets and their associated ontologies for describing different domains. According to LD principles, developers should reuse as many available terms as possible to describe their data. Importing ontologies or referring to their terms' URIs are the two main ways to reuse knowledge from available ontologies. In this paper, we have analyzed 18589 terms appearing within 196 ontologies included in the Linked Open Vocabularies (LOV) registry with the aim of understanding the current state of ontology reuse in the LD context. In order to characterize the landscape of ontology reuse in this context, we have extracted statistics about currently reused elements, calculated ratios for reuse, and drawn graphs about imports and references between ontologies.

**Keywords:** ontology, vocabulary, reuse, linked data, ontology import

## 1 Introduction

The Linked Data (LD) initiative enables the easy exposure, sharing, and connecting of data on the Web. Increasingly, datasets in different domains are being published according to the LD principles[1] and semantically well-defined using ontologies[2].

When particular data is going to be exposed as Linked Data, one of the first tasks [12] should be to develop an ontology that describes such data. Based on the guidelines for developing and publishing LD [7], the team involved in the project must develop such a vocabulary (a) reusing as many terms as possible from those existing in the vocabularies already published and (b) creating new proprietary terms, when available vocabularies do not model all the data that must be represented. Available vocabularies can be found in the Linked Open Vocabularies (LOV) registry[3].

The reuse of knowledge from available ontologies can be done in at least two different ways: (1) importing the ontology; and (2) referring to element URIs. In order to understand which is the current tendency of vocabulary reuse in the LD initiative, we consider it useful to analyse in depth how available vocabularies are reusing terms or (ontology) elements in the LD context. For this purpose, we have analyzed 73.96% of

1 http://www.w3.org/DesignIssues/LinkedData.html
2 For the purpose of this paper the terms "ontology" and "vocabulary" will be used indistinctly (http://www.w3.org/standards/semanticweb/ontology).
3 http://labs.mondeca.com/dataset/lov/index.html

adfa, p. 1, 2011.

the ontologies included in LOV. Vocabulary reuse can be considered as a kind of relatedness amongst ontologies [3]. In this paper, we characterize two different aspects of this reuse relatedness via an empirical study of LD vocabularies. These aspects are (a) the reuse by means of `owl:imports` statements and (b) the reuse by referring to element URIs.

To perform the analysis of these reuse aspects over the 196 ontologies, we have collected a set of static statistics. As a result of our study, we have (a) derived interesting metrics with respect to the current element reuse status and (b) sketched an overview of these reuse aspects in Linked Data vocabularies.

The rest of the paper is organized as follows: the experimental method followed to carry out the study is presented in Section 2. Section 3 shows the obtained results for the different tasks involved in the method as well as the analysis of such results. Finally, Section 4 exposes related research efforts while Section 5 presents some concluding remarks and future lines of work.

## 2 Experimental Method

The main aim of this study is to observe the current situation regarding vocabulary reuse in LD. Our study is only focused on the vocabularies registered in LOV[4, 5]; therefore, other ontologies used in any LD dataset are out of the scope of this work.

The method we followed in this experimental study is shown in Fig. 1. First, we harvested the vocabularies, second we extracted some static statistics and finally we calculated some derived statistics as well as drawing two different graphs.

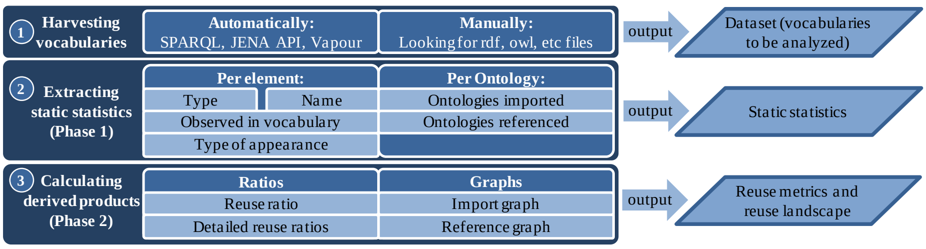


**Fig. 1.** Workflow to perform the study of reuse aspects over LD vocabularies

### 2.1 Harvesting vocabularies

In order to effectively harvest as many vocabularies as possible with minimum cost, we followed a semi-automatic approach. First, a SPARQL query was executed against the LOV SPARQL Endpoint to obtain the vocabularies and their preferred namespace prefixes.

For the obtained vocabularies, we have recorded the final URI (if any) retrieved by the server when requesting for “RDF/XML” content type within an http connection.

---

4 At the moment of performing the analysis, that is 12th of June, 2012.

5 Throughout this document those vocabularies gathered in the LOV registry will be named by their *vann:preferredNamespacePrefix* value defined in the LOV dataset.

For doing so, we have used the VAPOUR[6] API for Java. For those cases where there was no final URI retrieved by VAPOUR, we use the URI provided by LOV for downloading the vocabularies. After this process, we have filtered the obtained files using the JENA API[7]. We first tried to load each file into a JENA model and discarded those that either could not be loaded or they result into an empty model[8].

For those vocabularies that we were not able to download a file containing the RDF code or the downloaded file were not loaded correctly into a JENA model, we manually look for the ontology file using the URI provided by LOV registry as seed URI. These files manually downloaded were again loaded into an OWL JENA model. The final set of vocabularies taken into account in this study was composed by those that were successfully processed by JENA.

### 2.2 Extracting statistics

This process has been divided into two automated phases. During the first phase, we scanned the collected ontologies in order to track some statistics about those elements observed in each ontology as well as some statistic related to the ontology itself. It should be mentioned that in the context of this study we say that an element is "*observed*" in a given ontology if it appears as class, object property or datatype property in the JENA model after loading the ontology. During the second phase, we derived some statistics and defined two graphs from the recorded information.

**Phase 1: Static statistics for ontology elements and ontologies.** For each observed element in the harvested vocabularies we have recorded the following static information: (a) type of element (class, object property or datatype property); (b) element identifier (URI); (c) vocabulary in which the observed element (either locally defined or reused in any fashion) appears. This field contains the vocabulary prefix assigned in LOV; and (d) whether the element is *Local*, *Imported*, *Referenced* or *ReferencedByImport* in the vocabulary where it has been observed according to the classification shown in Fig. 2. This figure also includes an explanatory example together with the statements we can deduce for the example according to the abovementioned definitions. In addition, for each ontology we recorded the following static information: (a) imported ontologies; (b) original ontologies with elements referenced; and (c) original ontologies with elements referenced by means of (at least) one intermediate `owl:imports` statement. We consider this type of reuse as a more specific case of the reused by referencing elements.

**Phase 2: Derived statistics for ontology elements.** The aim of the second phase is twofold. On the one hand, we provide some ratios regarding different reuse metrics observed within the ontologies analyzed. On the other hand, we sketched a global view of the cloud(s) of vocabularies used in LD. In order to provide some measurements about the different types of reuse that we can observe in the ontologies, we have defined the following ratios extending the work presented in [13]. The *ReuseRa-*

---

6 http://vapour.sourceforge.net/api/

7 http://jena.sourceforge.net/

8 It should be mentioned that we only consider the ontology elements properly defined in RDF and therefore loaded into the JENA model.

*tio* represents the proportion of reused elements among all the elements appearing in a given ontology. The *ImportRatio* represents the proportion of imported elements, including the entire imports' closure, among all the reused ones. The *ReferenceRatio* represents the proportion of referenced elements among all the reused elements in a given ontology. Finally, the *ReferenceByImportRatio* represents the proportion of referenced by import elements among all the reused elements in a given ontology. Table 1 shows how these ratios are calculated.

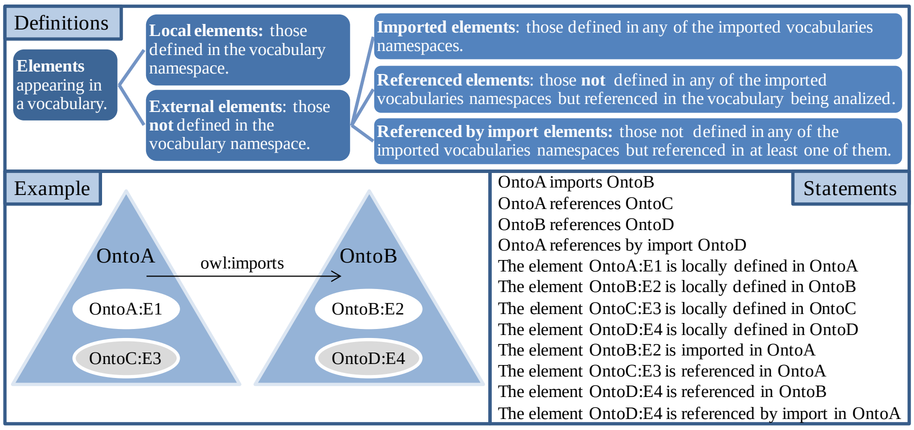


**Fig. 2.** Different roles played by elements in vocabularies and explanatory example

In addition, we have generated the following graphs from the static statistic we gathered for the ontologies: (a) the *ImportGraph* shows directed edges from any ontology with at least one `owl:imports` statement to the ontologies imported, and (b) the *ReferenceGraph* shows directed edges from any ontology with at least one referenced (or referenced by import) element to the ontologies from where the referenced elements were locally defined.

**Table 1.** ReuseRatio, ImportRatio, ReferenceRatio and ReferencedByImportRatio formulas

| | | |
|---|---|---|
| $ReuseRatio = \frac{\sum ReusedElements}{\sum TotalElements}$ | $ImportRatio = \frac{\sum ImportedElements}{\sum ReusedElements}$ | $ReferenceRatio = \frac{\sum ReferencedElements}{\sum ReusedElements}$ |
| $ReferenceByImportRatio = \frac{\sum ReferencedByImportElements}{\sum ReusedElements}$ | | |

## 3 Results, Analysis, and Discussion

Here we report, synthesize, analyse and discuss the results obtained[9] (a) for the process of harvesting vocabularies (Section 3.1), (b) for extracting static statistics (Section 3.2); and (c) for extracting derived statistic and graphs (Section 3.3). Finally, the results are discussed in Section 3.4.

[9] The set of 196 vocabularies analyzed, the obtained results (.csv files), and obtained graphs are available at http://www.oeg-upm.net/files/mpoveda/OEDW2012.

### 3.1 Findings for the process of harvesting vocabularies

Through our semi-automated method for harvesting vocabularies we have obtained a total of 196 vocabularies within the 265 registered in LOV. More specifically, we first retrieved 265 vocabulary URIs and their prefixes from the LOV SPARQL endpoint. By means of an http request (see Section 2.1) we obtained 242 files from which 190 were successfully loaded into a JENA model.

For those vocabularies that we were not able to download a file containing the RDF code (23 vocabularies) or to load it correctly into a JENA model (52 vocabularies), we manually looked for the file describing the ontology. After that, 56 files from which only 6 were successfully loaded into a JENA model. In summary we obtained a subset of **196** files, from which 190 were obtained automatically and 6 manually.

From this process we have realized that (a) some ontologies are difficult to find even when manually looking for the files containing the ontology in a given documentation website (e.g., "lsc" and "teach") and (b) others are not reachable due to connection problems with the host server (e.g., "adms" and "cgov"). In conclusion, in order to enhance the ontology reuse in LD, vocabulary publishers should ease (a) the tasks of accessing and processing vocabularies programmatically by implementing recommended methods for publishing vocabularies [1, 14] and (b) the task of finding and understanding the vocabularies for other developers by providing user friendly web sites where both the ontology and its documentation are easily accessible.

### 3.2 Findings for the process of extracting static statistics

In this section we provide the results and analysis of Phase 1.

**Static statistics for ontology elements**: The ontologies analyzed contain 18589 element appearances in total, covering a lot of real world domains, e.g., time, geography, life sciences, government, etc. and also general, upper level and metadata definitions. These element appearances are classified in Table 2 under two perspectives: (a) by their type of appearance, i.e., if they are locally defined in the vocabulary being analyzed, imported, referenced or referenced by an import and (b) by their type of element, i.e., if they are classes, object properties or datatype properties (See footnote 8).

Note that the same element can have different appearances, one in the vocabulary where it is locally defined and several in the vocabularies where it is imported or referenced. Concretely, there are 12032 different elements[10] appearing and the remaining 6557 are repetitions of some of those elements along the different vocabularies.

As the data in Table 2 indicates, 59.47% (11054 out of 18589) of the observed elements correspond to original definitions while the remaining 40.53% (7535 out of 18589) are consequences of reusing processes. More specifically, within these reused elements, 67.05% (5052 out of 7535) are reused by importing the original ontology and just 18.09% (1363 out of 7535) and 14.86% (1120 out of 7535) are reused by referencing to elements URI and by importing an ontology that references to ontology elements respectively.

[10] The analyzed vocabularies are those registered in LOV; however, they reuse elements defined either by vocabularies registered and not registered in LOV. For this reason, there are more elements appearing (12032) than defined within the analyzed vocabularies (11054).

**Table 2.** Classification of elements by type of element and by role played in the appearance

| | Classes | Object Properties | Datatype Properties | Total |
|---|---|---|---|---|
| Locally Defined | 5384 | 3956 | 1714 | 11054 |
| Imported | 1671 | 2297 | 1084 | 5052 |
| Referenced | 783 | 314 | 266 | 1363 |
| ReferencedByImport | 488 | 484 | 148 | 1120 |
| Total | 8326 | 7051 | 3212 | 18589 |

From the statistical data shown in Table 2, we can observe that 67.05% of the reused elements have been reused by importing ontologies while ontly 32.95% have been reused by referring to their URIs.

**Static statistics for ontologies**: On the one hand, one of the aims of our experiment is to know how many ontologies import other vocabularies among those used in LD. Once we gathered the 196 vocabularies we realized that there were 165 `owl:imports` statements within a total of 68 ontologies; that is, more than the 34% of the vocabularies analyzed use the `owl:imports` statement. Table 3 shows the 11 most popular imported ontologies and how many times they are imported. These popular imported ontologies in general are well documented and maintained, which is a possible reasons for their popularity.

**Table 3.** Most popular imported ontologies

| Imported ontology | Prefix | Times being imported |
|---|---|---|
| http://purl.org/dc/elements/1.1/ | dce | 15 |
| http://www.w3.org/2003/06/sw-vocab-status/ns | vs | 10 |
| http://purl.org/dc/terms/ | dc | 9 |
| http://xmlns.com/foaf/0.1/ | foaf | 9 |
| http://purl.org/NET/c4dm/event.owl | event | 8 |
| http://purl.org/goodrelations/v1 | gr | 5 |
| http://www.w3.org/2006/time | time | 5 |
| http://purl.org/vocab/vann/ | vann | 4 |
| http://purl.org/NET/scovo | scovo | 3 |
| http://purl.org/ontology/ao/core | ao | 3 |
| http://purl.org/ontology/similarity/ | sim | 3 |
| http://www.linkedmodel.org/schema/vaem | vaem | 3 |

On the other hand, this study also aims to know how many ontologies reference elements from other vocabularies. Among the 196 vocabularies analyzed, we observed that a total of 104 ontologies, more than the 53%, reference to other vocabularies. Table 4 shows the 11 most popular referenced ontologies and how many times they are referenced within the analyzed vocabularies.

### 3.3 Findings for the process of extracting derived statistics

In this section we provide the results and analysis of Phase 2.

**Derived statistics**: One of the aims of our experiment consists on calculating the reuse level in the current LD vocabularies. For doing so, we have calculated for each analyzed ontology its reuse ratio. In order to analyze in more detail the different types of reuse that we can observe in a given ontology we have calculated as well the reuse ratio due to (a) the use of the `owl:imports` statements, (b) references to other ontology elements and (c) references to other ontology elements by means of a `owl:imports` statement. Fig. 3 shows frequency distribution graphics for each of the ratios presented in Section 2.2. In these graphics the X axis shows the percentage (values between 0.0 and 1.0) of each type of reuse while the Y axis shows the number of ontologies that present a given percentage of reuse. Looking at the *ReuseRatio* bars we can see that a total of 101 ontologies present a reuse percentage between 0.0 and 0.1, that means that most of the ontologies do little or no reuse. In fact, the median for this distribution is 0.0 (value for the 98$^{th}$ ontology ordered by reuse ratio).

Focusing in the ontologies that do reuse in any of the considered types (*ImportRatio*, *ReferenceRatio* and *ReferenceByImportRatio*), we can observe that the trend is to adopt a type of reuse for each ontology, that is, most of the reuse is either based on `owl:imports` statements or based on referencing element URIs, however it is scarce to find ontologies combining both types of reuse at the same level. One of the ontologies that is completely based on referenced elements from other ontologies is the case of "gnm". The aim of "gnm" ontology is to establish mappings between DBpedia and Geonames ontologies, therefore all the appearing classes and properties are defined in DBpedia and Geonames ontologies instead of in the "gnm" namespace. However, for those cases with a reuse ratio higher than 60% the tendency is to achieve this level by importing ontologies. Not surprisingly, it could be due to the `owl:imports` statements mechanism that include and its transitivity.

**Table 4**. Most popular referenced ontologies

| Reused ontology | Prefix | Times being referenced |
|---|---|---|
| http://xmlns.com/foaf/0.1/ | foaf | 43 |
| http://purl.org/dc/terms/ | dc | 26 |
| http://www.w3.org/2003/01/geo/wgs84_pos | geo | 25 |
| http://purl.org/dc/elements/1.1/ | dce | 14 |
| http://www.w3.org/2004/02/skos/core | skos | 14 |
| http://www.w3.org/2000/10/swap/pim/contact | con | 11 |
| http://schema.org/ | schema | 8 |
| http://purl.org/NET/c4dm/event.owl# | event | 7 |
| http://dbpedia.org/ontology/ | DBpedia*[11] | 5 |
| http://purl.org/ontology/bibo/ | bibo | 5 |
| http://purl.org/vocab/frbr/core# | frbr | 5 |

**Graphs about ontology reuse**: Another goal within this study is to sketch the *ImportGraph* and *ReferenceGraph* (See Section 2.2). The former represents ontologies connected by an `owl:imports` statement while the later shows which ontologies are, directly or by means of an `owl:imports` statement, referenced by others.

It is worth noting that some ontologies from those registered in LOV import or reference to ontologies not registered in LOV. This fact is shown in the graphs in two ways: (a) the ontologies registered in LOV are denoted by their prefix in the nodes while the others are denoted by their URI and (b) they are represented by blue and orange nodes respectively.

[11] Prefixes marked with an * in this table refer to ontologies that are not included in LOV.

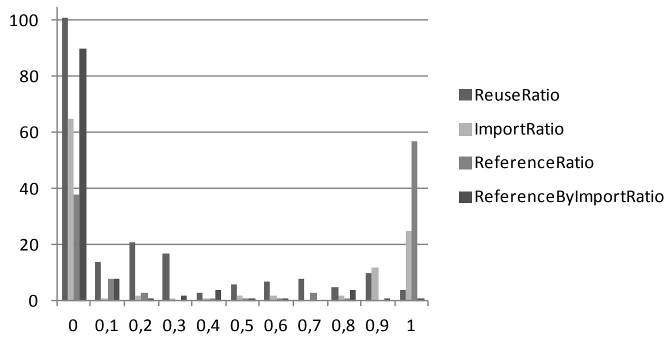


**Fig. 3**. Reuse ratios frequency distribution

Fig. 4 shows a bird's-eye view[12] of the *ImportGraph* and *ReferenceGraph*. It can be observed that both are unconnected graphs, the former with 28 connected components and the later with 3 connected components. Analyzing the *ImportGraph* in more detail we can observe that most of the ontologies act as a sink (only have in-links) (67) or as a source (only have out-links) (54) of import links while very few (14) have both in-links and out-links. For the *ReferenceGraph* 96 nodes are sink nodes, 79 are source and again 14 nodes have in-links and out-links. It should be noted that these graph characteristics, specially the high number of sink nodes, are due to the fact that the ontologies not registered in LOV are out of the scope of this study, therefore they are always sink nodes in the generated graphs.

### 3.4 Discussion

Looking closer at the figures presented in Sections 3.2 and 3.3, we can observe, on the one hand, that the reuse of ontology elements is applied in most of the cases (67.05%) by means of `owl:imports` statements while a minor part (32.95%) is due to element URI references. On the other hand, the tendency seems to be the reuse by referencing ontology element URIs as 104 ontologies do against 68 ontologies that import other ontologies. In addition, graphs in Fig. 4 (at the en of the paper) also reveal this as the *ReferenceGraph* is denser than the *ImportGraph*, which is rather sparse.

This apparently contradictory situation could be a side effect of the import mechanism and its transitivity, i.e., by stating a single `owl:imports` the whole content of another ontology is incorporated in the ontology importing it, as well as the complete imports' closure.

Finally, it is worth mentioning that ontology developers tend to reuse particular elements from other ontology by referencing them instead of reusing more terms than needed by means of `owl:imports` statements. This is an interesting fact as ontology editors support `owl:imports` through few simple user interactions while reusing part of an ontology involves more complex ontological engineering activities, for example: module extraction, partitioning, pruning, merging, etc.

[12] Detailed views of the graphs are available at http://www.oeg-upm.net/files/mpoveda/OEDW2012/Graphs/.

## 4 Related Work

Several research studies have been performed on available ontologies and datasets in LD. However, to the best of our knowledge, there are no specific analyses of the reuse relation among vocabularies, as we presented in this paper.

Although not directly related to LD, it is important to mention the work described in [15] about different statistical data extracted from available ontologies. This study was mainly focused on languages issues: the percentage of ontologies in the different OWL language species and the frequency of occurrences of OWL constructs.

Regarding the Web of Data, [4] presents an analysis of the use of OWL in linked data ontologies with respect to different features. Authors conclude that OWL is partially used in LD, and for this reason they propose a new profile, called OWL LD, more suitable for modelling linked data vocabularies.

With respect to studies over available datasets in LD, authors of [9, 8] have uncovered that many datasets are only loosely linked rather than tightly interconnected as recommended by LD principles [2, 7], and often only have links at instance level. There are also specific efforts for analyzing how LD datasets are related at the instance level. One of these efforts [6] included two types of datasets and analysed their relationship to one another, namely point-of-access datasets (e.g., DBpedia) and distributed datasets (e.g., the FOAF-o-sphere). Another interesting effort [5] presented the analysis of the (mis)-use of the `owl:sameAs` construct in LD.

Finally, regarding the relationship between vocabularies and datasets, [3] presented an empirical study of four different types of relationships between ontologies: (a) semantic relatedness, (b) content similarity, (c) expressivity closeness, and (d) distributional relatedness. In addition, in [10] authors described how RDF vocabularies are being used by the publisher in LD. In the line of elucidating how vocabularies are being used in the Web of Data, Nikolov and Motta's work [11] identified which schema-level relations can be extracted from existing data-level links.

## 5 Conclusions and Future Works

In this paper we draw the current reuse status in a subset of the vocabularies being used in LD, so that we are able to note several trends and make interesting observations. This study can be useful for different parties as (a) Linked Data working teams aiming to reuse ontology terms within their developments or (b) LOV developers to include new aspects and metrics of the vocabularies in their ecosystem.

We have first observed that the appearance of reused elements in the analyzed vocabularies is quite high (40.53%). However, more than a half (59.47%) corresponds to locally defined elements. Focusing on reused elements, we have seen that 67.05% are reused by importing ontologies, 18.09% are reused by referencing to elements URIs and 14.86% are reused by importing ontologies that reference to elements URIs.

In addition, we have sketched a first version of the linked vocabularies cloud overview. In this regard, our main line of future work is to complete the set of vocabularies analyzed so that all vocabularies appearing in the nodes are included in the study.

Finally, a future line of work is to analyze the outliers obtained from our study as some results might be due to mismatches between URIs (e.g., mismatch between a

URI used in an `owl:imports` statement and the one use as preferred in the ontology being imported) or mismatches between ontology versions (e.g., the ontology retrieved when importing a given namespace and the one found following an ontology documentation website).

**Acknowledgments.** This work has been partially supported by the Spanish projects *BabelData* (TIN2010-17550) and *BuscaMedia* (CENIT 2009-1026).

## References

1. Berrueta, D. Fernández, S., Frade, I. *Cooking HTTP content negotiation with Vapour*. Workshop on Scripting for the Semantic Web SFSW2008 Tenerife Spain. (2008)
2. Bizer, C., Heath, T., Berners-Lee, T.: Linked data - the story so far. Int. J. Semantic Web Inf. Syst 5(3), pp 1-22 (2009)
3. Cheng, G., Gong, S., Qu, Y. *An empirical study of vocabulary relatedness and its application to recommender systems.* In Proceedings of the 10th international conference on The semantic web - Volume Part I (ISWC2011), Springer-Verlag, Berlin, Heidelberg,. (2011)
4. Glimm, B., Hogan, A., Krötzsch, M., Polleres, A. *OWL: Yet to arrive on the web of data?* In WWW2012 Workshop on Linked Data on the Web (LDOW2012), Lyon, France, April 2012. (2012)
5. Halpin, H., Hayes, P.J., McCusker, J.P., McGuinness, D.L., Thompson, H.S. *When owl:sameAs isn't the Same: An Analysis of Identity in Linked Data*. In Proceedings of the 9th International Semantic Web Conference (ISWC2010). Shanghai, China. (2010)
6. Hausenblas, M., Halb, W., Raimond, Y., Heath, T. *What is the size of the Semantic Web*. Proceedings of ISemantics 2008, 9–16. (2008)
7. Heath, T., Bizer, C.: *Linked data: Evolving the Web into a global data space* (1st edition). Morgan & Claypool (2011)
8. Hogan, A., Harth, A., Passant, A., Decker, S. & Polleres, A. *Weaving the Pedantic Web*. Linked Data on the Web Workshop LDOW2010 at WWW2010 (2010)
9. Jain, P., Hitzler, P., Yeh, P.Z., Verma, K., Sheth, A.P.: *Linked data is merely more data*. In: AAAI Spring Symposium Linked Data Meets Artificial Intelligence. pp. 82-86. (2009)
10. Lorey, J., Abedjan, Z., Naumann, F., Böhm, C, *RDF Ontology (Re-)Engineering through Large-scale Data Mining*, Billion Triple Challenge, ISWC2011. (2011)
11. Nikolov, A., Motta, E. *Capturing emerging relations between schema ontologies on the Web of Data*. 9th International Semantic Web Conference (ISWC2010), Workshop on Consuming Linked Data (COLD 2010). Shanghai, China. (2010)
12. Villazón-Terrazas, B., Vilches-Blázquez, L., Corcho, O., Gómez-Pérez, *A. Methodological Guidelines for Publishing Government Linked Data Linking Government Data*. In Linking Government Data, chapter 2, pages 27–49. Springer New York, New York, NY, 2011.
13. Vrandečić, D. Ontology Evaluation. PhD Thesis. Karlsruher Instituts für Technologie. Germany. June 2010. (2010)
14. W3C Working Group Note. *Best Practice Recipes for Publishing RDF Vocabularies*. 28 August 2008. http://www.w3.org/TR/swbp-vocab-pub/ (2008)
15. Wang, T. D., Parsia, B., Hendler, J.A. *A Survey of the Web Ontology Landscape*. The Semantic Web - ISWC 2006, 5th International Semantic Web Conference. Athens, GA, USA. Pages: 682-694. ISBN: 3-540-49029-9. (2006)

## Annex

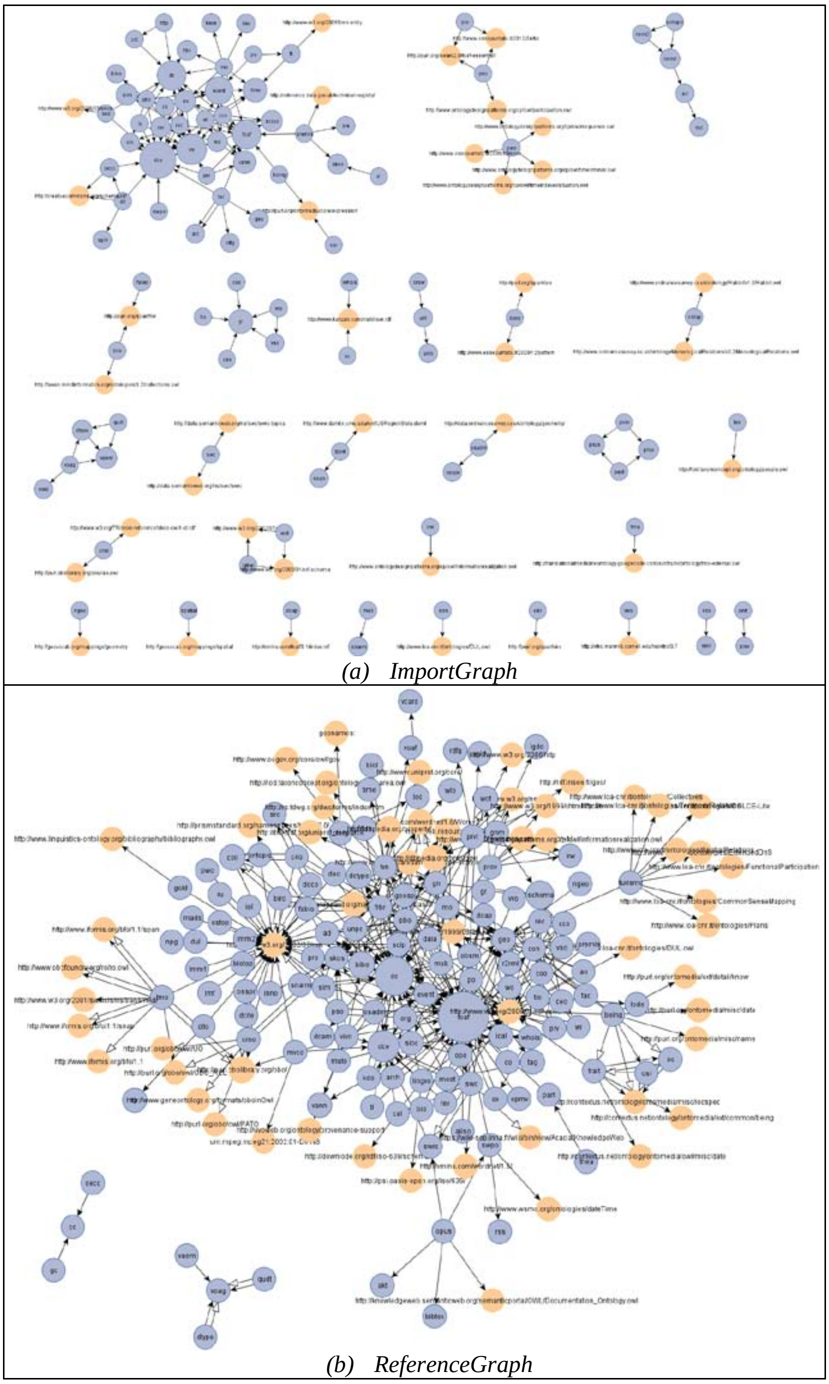

*(a) ImportGraph*

*(b) ReferenceGraph*

**Fig. 4.** *ImportGraph* and *ReferenceGraph* overview