Analyzing the Evolution of Linked Vocabularies

Mohammad Abdel-Qader[1,2][0000−0002−9702−1953], Iacopo Vagliano[2][0000−0002−3066−9464], and Ansgar Scherp[3][0000−0002−2653−9245]

[1] Christian-Albrechts University, Kiel, Germany
stu120798@mail.uni-kiel.de
[2] ZBW – Leibniz Information Centre for Economics, Kiel, Germany
{m.abdel-qader,i.vagliano}@zbw.eu
[3] University of Essex, United Kingdom
ansgar.scherp@essex.ac.uk

Abstract. Reusing terms results in a Network of Linked vOcabularies (NeLO), where the nodes are the vocabularies that use at least one term from some other vocabulary and thus depend on each other. These dependencies become a problem when vocabularies in the network change, e. g., when terms are deprecated or deleted. In these cases, all dependent vocabularies in the network need to be updated. So far, there has been no study that analyzes vocabulary changes in NeLO over time. To address this shortcoming, we compute the state of NeLO from the available versions of the vocabularies over 17 years. We analyze static parameters of NeLO such as its size, density, average degree, and the most important vocabularies at certain points in time. We further investigate how NeLO changes over time. Specifically, we measure the impact of a change in one vocabulary to others, how the reuse of terms changes, and the importance of vocabularies changes. Our analyses provide for the first time in-depth insights into the structure and evolution of NeLO. This study helps ontology engineers to identify shortcomings of the data modeling and to assess the dependencies implied with reusing a specific vocabulary.

1 Introduction

For modeling and publishing data on the web, we use properties and types defined in one or multiple vocabularies. It is common practice to reuse existing terms, i. e., properties and types, from other vocabularies for modeling one's own data. The goal is to prevent the proliferation of terms and to reduce the range of choices when modeling data. This reuse of terms leads to a Network of Linked vOcabularies (NeLO). In essence, NeLO is a directed graph of connected vocabularies that have at least one reuse from some other vocabulary. By connected vocabularies, we mean that a vocabulary v is reusing at least one term from another vocabulary w.

The connections between the vocabularies become a problem when one or more of the vocabularies in the network change. For instance, the vocabulary w could declare a term t as deprecated or even delete it while the dependent vocabulary v is reusing this term t. The changes of vocabularies have a direct impact on all dependent vocabularies, i. e., those that reuse any of the changed terms. Furthermore, all the data that are modeled with these outdated vocabularies have also to be updated. The outdated terms are those that were deleted or deprecated when updating a vocabulary.

Previous research focused on analyzing the interlink at an instance level. In contrast, with analyzing NeLO, we focus on the evolution of the web of data at the schema level. In a previous work [1], we showed that some deleted and deprecated terms are still reused by data publishers to represent their data. In this paper, we consider the reuse of vocabulary terms in other vocabularies. Specifically, we analyze NeLO by addressing the following research questions:

RQ1 What is the state of the Network of Linked Vocabularies? This includes several subquestions: What is its size in terms of the number of nodes and edges? What is its density, and average degree? Which are the important vocabularies, i. e., central nodes?

RQ2 How are vocabulary terms reused by other vocabularies? More specific subquestions are: How many vocabularies do reuse terms from others? How many terms are reused? Are the reused terms the most recent ones? How does the change (addition or deletion) of terms in one vocabulary impact the other vocabularies on the network?

RQ3 How do ranking metrics, such as *PageRank*, *Hypertext Induced Topic Selection* (HITS), and *Centrality*, change during the evolution of the Network of Linked Vocabularies? We are specifically interested in understanding how the important nodes, i. e., the central vocabularies, as well as the reuse of terms changes over time.

To address these questions, we analyzed 994 vocabularies and their changes in a time span of over 17 years. We considered vocabularies as part of the network if they import or export at least one term from some other vocabularies. We employed a broad range of network-analysis metrics on the extracted network and applied them during the evolution of NeLO to find out how the important nodes change over time. We investigated how the change of one vocabulary impacts the others that reuse its terms.

Our analysis shows that at the beginning the growth of the Network of Linked Vocabularies was large, but recently the increase has been lower. Moreover, the percentage of reused terms by other vocabularies has decreased over time. This study also summarizes how the reused vocabularies changed over time. Overall, we believe that our study can help ontology engineers by raising awareness on the changes occurred in NeLO. This may lead to an increase of the reuse of terms among vocabularies, and to avoid or decrease terms' redundancy.

The remainder is structured as follows. We introduce a motivating example for analyzing NeLO in Section 2. We review related work in Section 3. In Section 4, we describe our experimental apparatus to analyze the evolution of NeLO. We present our results in Section 5 and discuss them in Section 6, before we conclude.

2 Motivating Example

Figure 1 shows a selected part of NeLO, with some of the dependencies of the vocabularies depicted. The arrows represent the relation between exporters and importers. An arrow from a vocabulary w to another vocabulary v indicates that v imports terms from w, or, in other words, that w exports terms to v. The size of the nodes represents the number of exports, i. e., more exports imply a bigger node. The width of the edges represents the total number of types and properties that the target vocabulary imports from the source vocabulary. For example, the *adms* vocabulary exports terms to *food*, *gn*, *search*, and *void*, while *schema*, and *voaf* export terms to *adms*.

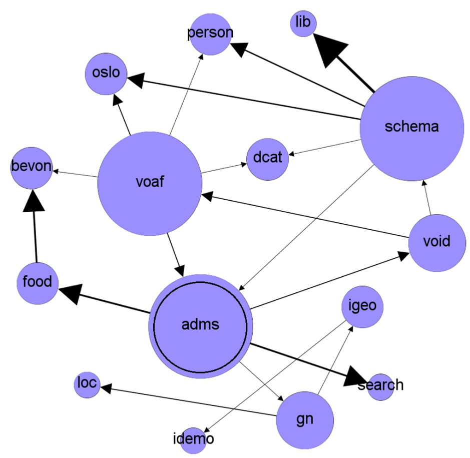

Fig. 1: Vocabularies that import terms from *adms* and other vocabularies.

Regarding the evolution of NeLO, we consider the example of the *adms* vocabulary. *adms* deals with describing highly reusable metadata and reference data, which are called Semantic Assets[4]. Figure 2 shows the evolution of the *adms* vocabulary within six versions over five years (bottom). Furthermore, the *food* vocabulary (top) reuses some terms of *adms* and also has different versions over its lifespan. The *adms* vocabulary published six versions between May 2012 and July 2015, and introduced the `adms:SemanticAsset` type and `adms:accessURL` property in its version published in June 2012 (V2). The *food* vocabulary reuses those two terms in its first version, which were published in November 2012. Afterwards, a new version of the *adms* vocabulary has been released in May 2013, which deleted the `adms:SemanticAsset` and `adms:accessURL` terms. In September 2013, the *food* vocabulary was updated, but the updated version of the *food* ontology kept using the two terms that were deleted from *adms*.

[4] `https://www.w3.org/TR/vocab-adms/`

Such a scenario may mean that *food* still needs the deleted terms and its ontology engineers have found no alternatives. However, it could also denote that the ontology engineers of the *food* vocabulary are not aware of the changes in *adms*. This study analyzes the problem of evolution in NeLO and the possible effects that changes in vocabularies have on the dependent vocabularies.

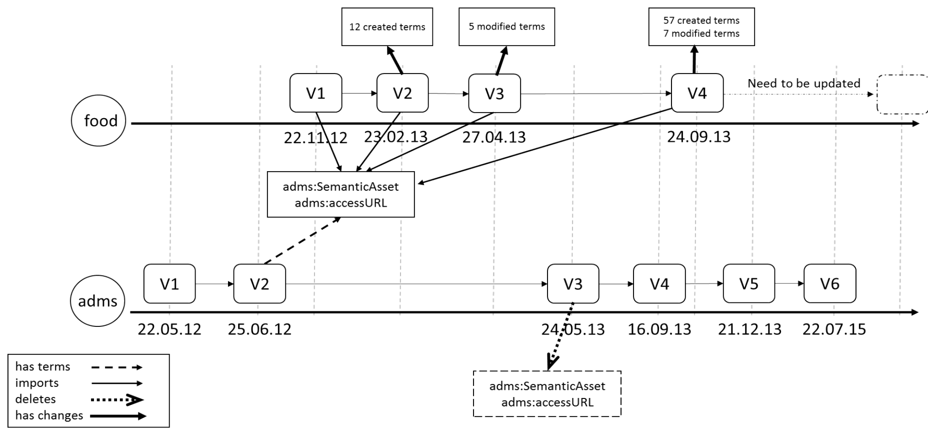

Fig. 2: Evolution of *adms* and its relation to its importer *food*.

3 Related Work

Several works analyzed the reusability and evolution of vocabularies. Some of these focused on biomedical ontologies. Reis et al. [4] studied how changes of ontologies affect the mappings between them. The goal was to understand the evolution of biomedical ontologies to propose an automatic mechanism for mapping. Hartung et al. [7] selected 16 life science ontologies since 2004 to measure the impact of ontology evolution on semantic annotations. They proposed a framework to analyze the life science ontologies and their instances. Cardoso et al. [2] analyzed the impact of ontology evolution on existing annotations. They considered over 66 million annotations from 5,000 biomedical articles and ontologies to support semi-automatic annotation maintenance mechanisms. Ghazvinian et al. [5] studied the overlap between the biomedical ontologies. They found more than 4 million mappings between concepts. Using those mappings, they analyzed the ontologies, their repositories, and how they can help in ontology design and evaluation. Kamdar et al. [11] published a study regarding terms reuse on ontologies in the BioPortal repository. The authors found reuse between 25-31%, and the percentage of reused terms was less than 9%. However, none of these studies applied network analysis metrics to the evolved ontologies. Furthermore, they studied the mappings and overlap between ontologies in the biomedical domain, while we analyze the vocabularies from various domains.

Vandenbussche et al. [18] described Linked Open Vocabularies (LOV) and provided some related statistics. They also provided a system that shows the dependencies between vocabularies, but it does not give information about which terms are reused by other vocabularies. In contrast, we provide information about the reusability of terms in NeLO, such as the most reused terms and whether terms, which have been deleted, are still reused or not. Most of the analyses of Linked Open Data focused on the instance level. Vassilis et al. [16] discussed the state-of-the-art systems that manage the evolving RDF data by proposing a benchmark generator that evaluate the ability of the current versioning strategies to manage LOD datasets. Käfer et al. [10] collected 29 weekly snapshots of a seed list with 86,696 RDF documents and analyzed the changes between pairs of two consecutive snapshots. The results showed that RDF documents change frequently. Gottron and Gottron [6] compared the accuracy of various RDF indices over the weekly snapshots from Käfer et al. [10]. Dividino et al. [3] analyzed the dynamics of the data by Käfer et al. [10] and proposed a monotone, non-negative function to represent the dynamics of RDF statements as a single numerical value. Nishioka and Scherp [13] computed periodicities of temporal changes in the dataset by Käfer et al. [10].

Palma et al. [15] proposed guidelines to execute the ontology evolution activity. Their approach covered two aspects: the description of the ontology evolution process and the tasks involved, and the facilitation of the process using semi-automatic techniques. While their methodology is to undertake the evolution process, we analyze the evolution in NeLO. Meusel et al. [12] analyzed the evolution of *schema.org* over four years. Thus, they focused on analyzing only a single but widely used data schema. They studied the top-down adoption and bottom-up evolution approaches and found that some of the deprecated terms are still used. Noura et al. [14] identified the most popular ontologies on the Internet of Things to identify the most used terms in this domain. They selected 14 ontologies. They found out that 71% of the ontologies reuse less than 18% of the terms defined, and 20% of ontologies are not reused at all. Jiménez-Ruiz et al. [9] described a logic-based approach to reuse terms between ontologies. Their approach specified that the reuse should be safe, i. e., the reused terms are valid (have not been changed/deleted in the source). Furthermore, the reuse should be economic, i. e., only the relevant parts of an ontology are imported. Previously, we presented a qualitative assessment of vocabulary changes [1], but we focused on their impact on instances. In this paper, we complement our prior work and consider the schema-level, i. e., the reuse of vocabulary terms in other vocabularies and the impact of their changes on each other.

4 Network Analysis Method

To answer our research questions, we conducted the following steps. First, we extracted all types and properties from all the available versions of vocabularies (from June 2001 to June 2018), which are listed in the Linked Open Vocabularies (LOV) dataset[5]. The terms extracted are classified into two categories: the *own terms* are the terms created by the ontology engineers of the considered vocabulary, while the *reused terms* are the terms that are reused from other vocabularies. Second, we employed different network-analysis metrics to study the Network of Linked Vocabularies, such as degree, PageRank, and HITS. We checked if the *reused terms* are the most recent ones, i. e., whether the terms that appear in the latest published version of the source vocabulary are actually those that are reused in the target vocabulary. This procedure was repeated on the evolving NeLO in a yearly basis to analyze the change of the Network of Linked Vocabulary over time.

For the first step, we examined 636 vocabularies listed in LOV. We employed the OWL API[6] version 5.1.6 to extract all the *own terms* and *reused terms* from the latest version of all the 636 vocabularies. While extracting the reused terms, some additional vocabularies that are not contained in LOV, were found. Thus, we considered a total of 994 ontologies. For the second step of the methodology, we used the Open Graph Viz Platform (Gephi)[7] version 0.9.2 to visualize and analyze the Network of Linked Vocabularies. Subsequently, we identified the deleted and deprecated terms of the vocabularies by parsing and comparing all versions of a vocabulary. Finally, we checked if the vocabularies on NeLO are still reusing the deleted or deprecated terms.

5 Results

We present the results of our analysis in Sections 5.1 to 5.3 along the research questions RQ1, RQ2, and RQ3. We only present the main results. The complete results can be found online.[8]

5.1 State of NeLO in 2018

Figure 3 shows the current state of NeLO after extracting all import relations between the latest versions of the vocabularies until June 2018. One can see three main circles in the network. Those circles are formed depending on the number of exports to the other vocabularies. The inner circle contains the vocabularies that have the most exports (more than 100 edges), which are represented by the larger node sizes. The middle circle (the denser area of smaller nodes) includes the vocabularies which have between 5 and 100 edges. The outer circle (the sparser external area) contains all the vocabularies that have been imported by less than five vocabularies. A fully scalable version of the figure is available at our website.[8]

In June 2018, NeLO consists of 994 vocabularies and 7,046 edges between those vocabularies, with a density of 0.007. Thus, the actual number of edges in

[5] `http://lov.okfn.org/dataset/lov/`, last accessed: June 2018
[6] `https://github.com/owlcs/owlapi`, last accessed: June 2018
[7] `https://gephi.org/`, last accessed: June 2018
[8] `https://sites.google.com/view/nelo-evolution`

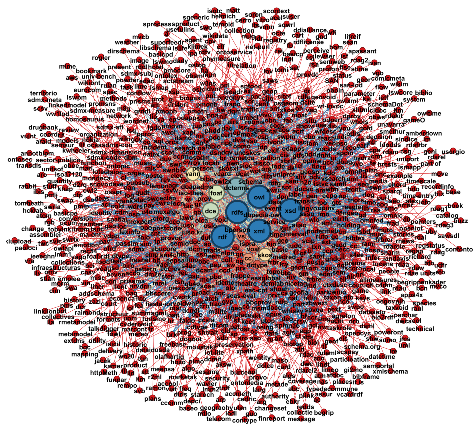

Fig. 3: The Network of Linked Ontologies (NeLO) as it appears in June 2018.

the graph is far away from the maximal number of possible edges (when each node has an edge to all other nodes) and the maximal density (equal to 1 with the maximal number of possible edges). The average degree for NeLO 2018 is 7.09, with a standard deviation of 7.46.

Tables 1, 2, and 3 list the top-10 vocabularies that have the highest scores for degree, HITS, and PageRank, respectively.[9] PageRank and HITS can help to identify nodes which can be problematic because they have many dependencies, or their changes may affect many other nodes because they are widely reused. We exploit these measures in addition to the degree since they take into account indirect dependencies (indirect links in NeLO). Most of the vocabularies are the same for all these metrics, with some differences in their order. Furthermore, *dcterms*, *dce*, *foaf*, *skos*, and *vann* appear in the top of the three tables.

Please note, we remove meta-vocabularies since it is quite natural that they are mostly used. Since we could not find a clear definition of meta-vocabularies, we excluded *owl*, *rdf*, *rdfs*, *xml*, and *xsd*, which clearly belong to this category. Some other vocabularies may be considered meta-vocabularies, too. For example, *cc* and *vann* are normally used to annotate metadata of the ontology itself. Our data are publicly available,[8] and we welcome researchers to recompute the results excluding these or other ontologies.

5.2 Reuse of Vocabularies and Adoption of their Changes

Reusing existing terms is one of the main principles of Linked Data. Table 4 lists the top-10 terms reused by other vocabularies. Those terms are extracted from the latest NeLO snapshot (June 2018). After excluding the meta-vocabularies, the most reused term is `dcterms:modified`, which represents the date on which

[9] We refer to Zaki et al. [19] for a description of degree, HITS, and PageRank.

Table 1: Top-10 vocabularies for Degree, In-degree, and Out-degree in 2018, sorted by Degree. The scores are calculated over both types and properties.

Vocabulary	Degree	In-degree	Out-degree
dcterms	435	425	10
dce	347	339	8
foaf	330	317	13
vann	255	244	11
skos	235	229	6
cc	153	146	7
voaf	121	103	18
vs	116	108	8
dctype	82	74	8
schema.org	73	61	12

Table 2: Top-10 vocabularies for HITS (Hub and Authority) scores in 2018, sorted by Authority.

Vocabulary	Authority	Hub
dcterms	0.305421	0.037978
dce	0.242374	0.037727
foaf	0.234664	0.044112
vann	0.184754	0.045030
skos	0.171827	0.034529
cc	0.113386	0.034723
vs	0.081972	0.040256
voaf	0.080739	0.045920
dctype	0.058152	0.037727
schema.org	0.046659	0.040364

Table 3: Top-10 vocabularies for PageRank in 2018.

Vocabulary	PageRank
dce	0.045954
dcterms	0.027649
skos	0.017678
foaf	0.013986
dcam	0.009152
vann	0.009117
grddl	0.008740
dctype	0.005744
cc	0.005446
vs	0.005005

Table 4: Top-10 terms that are reused by other vocabularies in 2018.

Term	Importing vocab.
dcterms:modified	281
dcterms:title	276
dce:title	266
dce:creator	263
vann:preferredNamespacePrefix	257
dcterms:description	249
vann:preferredNamespaceUri	241
foaf:Person	175
foaf:name	164
cc:license	122

a resource was changed. The term `dcterms:modified` has 281 vocabularies that reuse it.

Figure 4 shows a histogram of the vocabularies that reuse outdated terms from other vocabularies in the 2018 snapshot. We can notice that 16 vocabularies reuse one outdated term. On the other hand, we found six vocabularies that reuse more than six outdated terms.

There are three vocabularies that removed the reused terms after they were deleted from their original vocabularies, which are listed in Table 5. The *Updated version* column represents the version of the vocabulary where the update occurred. Notably, the *oslo* vocabulary removed five outdated terms, but it still reuses two outdated terms in its latest version.

5.3 Evolution of NeLO

Figure 5 shows the total number of available types and properties, and the total number of reused terms in NeLO. The reuse percentage was at its top with 10%

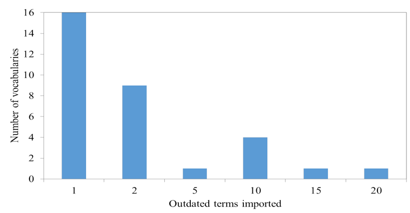

Fig. 4: The number of vocabularies that reuse outdated terms by the number of outdated terms reused.

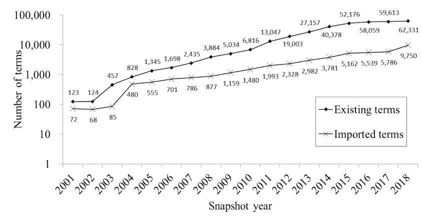

Fig. 5: The total number of existing types and properties and the terms reused from other vocabularies.

Table 5: Vocabularies that removed outdated terms.

Vocabulary	Removed terms	Updated version	Prior version
qudt	12	9-Oct-2016	1-Jun-2011
oslo	5	30-May-2014	30-Sep-2013
dcat	1	28-Nov-2013	20-Sep-2013

and 11% in 2010 and 2011, respectively, while for all other years it remains in the range between 5% and 7%.

Figure 6 depicts the total number of nodes and edges for each NeLO snapshot. It is worth noting that the number of nodes (vocabularies) and edges almost doubled from the 2003 to the 2004 snapshots compared to 2002 and 2003, respectively. Then they continued to roughly double every two years until 2013. After that year, the growing-rate decreased, and, since 2016 until June 2018, the number of new vocabularies that entered the network becomes small (around 70 new vocabularies per year), while the number of new links is still slightly higher (about 600 per year).

Figure 7 presents the density, network diameter [19], and average degree measures over time. The network average degree has a slow but steady increase. The density of the network is slightly decreasing over time. More specifically, in 2001, the network density was 0.273, and in 2018 it was 0.007. The network diameter sharply grew over the period considered, although its increase is not steady. First, it quadrupled from 2002 to 2003, then there is another small peak from 2004 to 2005. From 2010 to 2015 we can see the highest growth. The diameter of 2015 also represents the maximum value in the whole period. Finally, in the last three years, it is almost constant.

Figures 8 illustrates the evolution of the in-degree and out-degree metrics for the top-5 vocabularies, respectively. We selected the top-five vocabularies for these measures in the latest snapshots of NeLO (from 2015 to 2018), excluding the meta-vocabularies. Subsequently, we calculated the scores for those top-five vocabularies for each NeLO snapshot. This selection process holds also for the following analyses that consider PageRank and HITS (Figures 9 and 10).

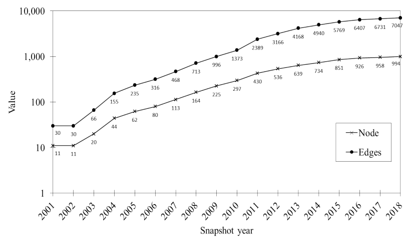

Fig. 6: Number of nodes and edges in NeLO.

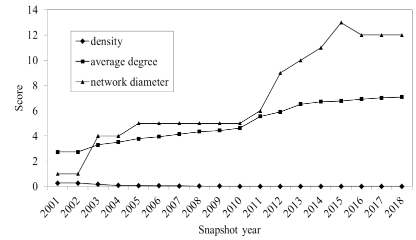

Fig. 7: Density, average degree, and network diameter.

Figure 8a depicts the number of imported vocabularies. While the out-degree for the vocabularies selected tends to steadily grow, the in-degree is mostly constant and abruptly increased in 2011 for *mo*, in 2015 for *interval*, and in 2018 for *semio*, with some exceptions. We can notice that *qudt* decreases the number of imported vocabularies. This number was 39 in 2011, then increased to 44 in 2012. Subsequently, it has continuously decreased to 25 imported vocabularies. Furthermore, the *oa* vocabulary decreases the number of imported vocabularies from 23 in 2013 to only 9 in 2016. Later, this number has increased again to reach 27 imported vocabularies. The *mo* vocabulary shows a constant number of imports from 2011 until 2018. While *mo* was introduced in 2007, it did not reuse any term from the other vocabularies until 2011.

Figure 8b presents the out-degrees for the top-5 vocabularies. The out-degree corresponds to the total number of other ontologies that reuse at least one term from those vocabularies, i. e. the number of exports to different ontologies. From 2003 to 2007, all the vocabularies shown have a similar out-degree. From 2009, *skos* started to increase more than the others, and the same holds for *vann* starting from 2012. We can notice that *vann* and *skos* have become widely more popular than the other vocabularies. Additionally, from 2015, *vann* exceeded *skos*, while earlier *skos* had the highest out-degree overall. However, the gap between their out-degrees is rather small. In 2014, *cc* achieved about the same out-degree of *vs*, and later on *cc* has a higher value than *vs*. The *voaf* vocabulary is introduced in 2011 and in 2018 accounts for almost the same in-degree as *vs*.

Figures 9 and 10 show the PageRank and HITS scores, respectively, for the same top-five vocabularies selected as for the degree analysis. In Figure 9, we can notice that all vocabularies have decreasing PageRank scores except *skos* and *vann*. The *skos* vocabulary started to increase its score from 2009, although from 2013 to 2018 it is again steady. However, this is almost half than the original *skos*'s PageRank score in 2003. Instead, *vann* had its lowest point in 2010, and started to slowly grow again from 2011. The *grddl* vocabulary appeared in 2008, with the lowest PageRank score, although it was close to *dctype* and *vann*. It slightly decreased in 2009. In 2010, it increased and remained almost constant in the following years, with roughly the same value as *dcam*.

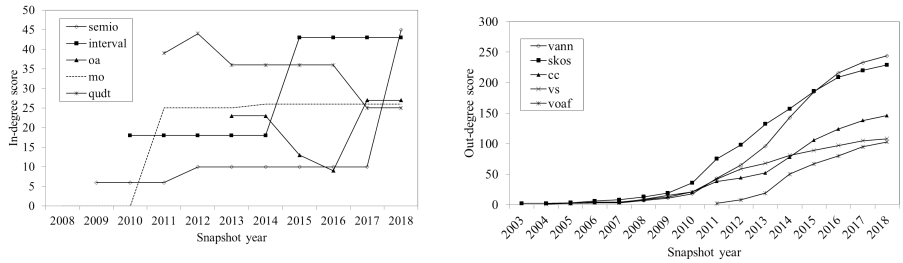

(a) The in-degree scores. (b) The out-degree scores.

Fig. 8: The in- and out-degree for the top-5 vocabularies on each NeLO snapshot.

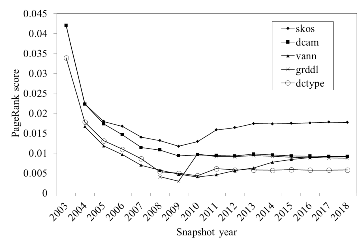

Fig. 9: The PageRank scores for the top-five vocabularies for each year.

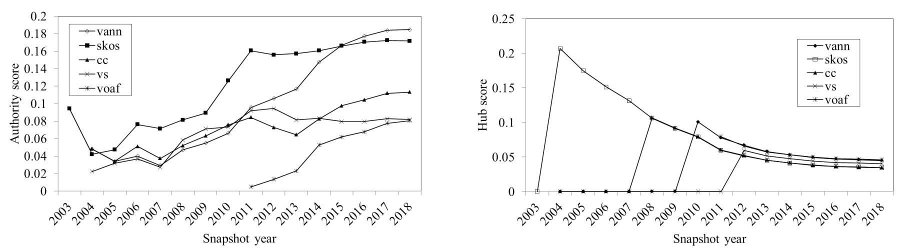

(a) The authority scores. (b) The hub scores.

Fig. 10: The HITS scores for the top-five vocabularies on each NeLO snapshot.

Regarding the HITS scores, Figure 10a shows that there is a general trend of increasing authority scores for all the vocabularies, although with some fluctuations. Specifically, *vann* started to grow from 2007, after an initial slight decrease. In 2018, it achieved the highest authority score. The *skos* vocabulary has a similar trend, with a more pronounced initial decrease from 2003 to 2004 and a peak in 2011. Subsequently, there is almost no further growth. Notably, *vs* has a score decrease starting from 2013, and then the score becomes stable. The *voaf* vocabulary appeared in 2011 and has steadily grown until 2018, where it achieved the same value as *vs*. The latter has the lowest scores among the vocabularies presented. Regarding the hub scores depicted in Figure 10b, the vocabularies show a similar pattern of a continuously decreasing score after an initial peak. The difference is in their peak value and in the year. Note that the early versions of vocabularies had no terms imported by other ontologies. Afterward, they started to be reused. In 2018, all the vocabularies achieved similar hub scores, around 0.05.

6 Discussion

6.1 State of NeLO in 2018

The vocabularies in NeLO form three categories (the three circles in Figure 3 introduced in Section 5.1). The first one corresponds the vocabularies, including the meta-vocabularies, that export terms to most of the other vocabularies in the network. These vocabularies in the central circle are the most important in NeLO 2018. They are the most popular in the sense that their terms are highly reused, but updating their terms is critical because of their potential impact on many other vocabularies which reuse their terms. Nevertheless, these vocabularies change rather rarely. In fact, they have on average three versions over 17 years. Overall, the vocabularies in this category represent 2% of all vocabularies, export their terms to 71% of the other vocabularies, and account for 66% of outgoing links. Vocabularies in the second category still have many edges to other ontologies, but less than the meta-vocabularies. These are also very popular, and updating them could impact various vocabularies. The average number of versions of the second category of vocabularies is around three. Thus, the vocabularies in this category seem to be more stable. The vocabularies in the middle circle account for around 20% of the outgoing links. These vocabularies represent 13% of the vocabularies, and their terms are reused by 56% of other vocabularies in NeLO 2018. The third category contains rarely-reused vocabularies, such as the newcomers, or the ones that cover a very specific domain.

6.2 Reuse of Vocabularies and Adoption of their Changes

Overall, 16% of the terms in NeLO are reused in June 2018. This number is still low and there is a need to increase the reuse of the existing types and properties, in order to avoid overlap and redundancy in the data representation [8]. Tools to suggest existing terms like TermPicker [17] could play a major role in increasing the number of terms reused by helping ontology engineers to select and discover terms to reuse.

Many vocabularies are up-to-date in NeLO 2018. There are 35 vocabularies that are affected by term updates in other vocabularies. 33 vocabularies are still using outdated types or properties. Although this number may seem low, it can have a strong impact on the published data, as shown in our previous study [1]. The number of outdated terms reused by those vocabularies ranges between 1 and 20. We think that the process of checking for changes in order to update the ontologies is done manually. The SemWeb Vocabulary Status ontology (*vs*)[10] provides information about the status of a term, but it is not widely used. This vocabulary (or similar ones) can help ontology engineers to check the recent status of terms before reusing them, e. g. to avoid reusing terms which are not stable and are likely to be removed in the future.

From the 35 vocabularies affected by changes in 2018, three have been updated by removing some of the outdated terms. For instance, the *oslo* vocabulary removed five terms, one from *adms* and four from *rov*. However, *oslo* still reuses two terms from *vcard*, although they have been deleted in *vcard*. This could either mean that the deleted terms are still needed and no alternatives have been found, or that some updates have been missed because the process for looking for changes in the other vocabularies is done manually. Reusing terms from older vocabulary versions, which can still be accessed by the IRI of the version, is possible, but we recommend checking the reason and update such terms.

There is a lack of tools to notify ontology engineers about changes in the vocabularies. Such tools may help ontology engineers to keep track of the changes and reduce the update effort. Tool support becomes especially important when a vocabulary has many dependencies: the more terms the ontology reuses, the higher is the effort to update the vocabulary when a change occurs. With many dependencies, it is challenging to keep an ontology up-to-date as any change in one of the imported vocabularies could require an update of the importer. Some vocabularies have edges from more than 40 others. Overall, 12% of the vocabularies imported from 59% of others, accounting for 22% of incoming links.

6.3 Evolution of NeLO

The number of new vocabularies and relations between them has decreased over time. While in 2003, 55% of the vocabularies were new, this percentage decreased to 4% in 2018. Regarding the edges, 57% of them were introduced in 2003. This percentage decreased to 27% in 2009, increased to 43% in 2010, and dropped to 4% in 2018. We can observe fluctuations in the number of new vocabularies and edges. Ontology engineers keep adding terms to their existing vocabularies, rather than introducing new ontologies, in order to fulfill their domain requirements. Therefore, over time we expect that the number of new vocabularies will continue to decrease or perhaps there will be a slower growth rate. Given that less new vocabularies have been introduced over time, it is not surprising that also less import/export links have been created.

Considering the reuse of terms from 2004 to 2010, the percentage of reused terms with respect to the available ones ranges between 58% and 22%. This percentage decreased to 10% in June 2017, although slightly increased in 2018, accounting for 16% of the available terms. This suggests that reusing terms was initially more common. One reason could be that initially much fewer vocabularies were available, it was easier to be aware of them and reuse their terms.

[10] `https://www.w3.org/2003/06/sw-vocab-status/note`

Nevertheless, more specific vocabularies, which are less suitable to be reused, may have been created over time.

Some vocabularies have become more popular (their out-degree has increased). When excluding the meta-vocabularies, *vann*, *skos*, *cc*, *vs*, and *voaf* are the most popular vocabularies. By taking into account the out-degree and centrality measures on all NeLO snapshots for the vocabularies with the highest scores in the last three years, we found that *vann*, *skos*, *cc*, *vs*, and *voaf* have increased their scores. Notably, *vann* and *skos* have a more rapid increase than the other three vocabularies, i. e., they have become more popular over time. Overall, the meta-vocabularies, which are suitable for most domains, are the most popular ones. Interestingly, our findings show a decline in the growth of out-degree scores, i. e., the average number of exports per vocabulary decreases over time. This could be due to the fact that less new vocabularies have been introduced over time. Consequently, fewer terms are exported to those new vocabularies. Nevertheless, the reuse of terms could still be increased among existing vocabularies, according to the needs of the particular application scenario considered. Regarding the in-degrees, we observed that they vary among the nodes in the network over time. Some of the vocabularies with the highest in-degree over time, such as *mo*, *interval*, and *semio*, have a sudden and large growth of imports at a specific point in time. This corresponds to a new version with a considerable extension of the previous vocabulary which reuses many terms from other ontologies. Thus, more effort is needed to keep track of the changes in the reused terms.

Similarly, the changes in the vocabularies with high PageRank and HITS scores affects many other vocabularies. The difference between those with high PageRank and HITS scores with a high out-degree is that their changes can significantly impact also ontologies that are indirectly related to it. Therefore, these changes can be even more critical. We recommend that ontology engineers of vocabularies that reuse terms from vocabularies with high PageRank and HITS scores periodically check them for changes.

7 Conclusion

By this study, we aim to raise ontology engineers' awareness about the changes in NeLO. As our analysis of the evolution of NeLO shows, the dynamics of changes has slowed down after some fast evolution between 2001 and 2010. As of today, 33 of the considered 994 vocabularies do reuse outdated terms. We recommend ontology engineers to check the evolution in NeLO and assess why a term is deleted or deprecated. Furthermore, we like to further stimulate an increase in reusing terms to prevent redundancy. As future work, we will consider the other types of updates, such as adding/removing constraints to terms or their subclasses, and different types of reuse of terms between vocabularies, e. g., introducing sub-classes or sub-properties, or using terms for annotation.

Acknowledgment This work was supported by the DFG (German Research Foundation) with the LOC-DB project (Grants No. GZ:SCHE 1687/5-1) and the EU's Horizon 2020 programme under grant agreement H2020-693092 MOVING.

References

1. Abdel-Qader, M., Scherp, A., Vagliano, I.: Analyzing the evolution of vocabulary terms and their impact on the LOD Cloud. In: ESWC. Springer (2018)
2. Cardoso, S.D., Pruski, C., Da Silveira, M., Lin, Y.C., Groß, A., Rahm, E., Reynaud-Delaître, C.: Leveraging the impact of ontology evolution on semantic annotations. In: European Knowledge Acquisition Workshop. pp. 68–82. Springer (2016)
3. Dividino, R., Gottron, T., Scherp, A.: Strategies for efficiently keeping local linked open data caches up-to-date. In: ISWC. pp. 356–373. Springer (2015)
4. Dos Reis, J.C., Pruski, C., Da Silveira, M., Reynaud-Delaître, C.: Understanding semantic mapping evolution by observing changes in biomedical ontologies. Journal of biomedical informatics **47**, 71–82 (2014)
5. Ghazvinian, A., Noy, N.F., Jonquet, C., Shah, N., Musen, M.A.: What four million mappings can tell you about two hundred ontologies. In: ISWC. pp. 229–242. Springer (2009)
6. Gottron, T., Gottron, C.: Perplexity of index models over evolving linked data. In: ESWC. pp. 161–175. Springer (2014)
7. Hartung, M., Kirsten, T., Rahm, E.: Analyzing the evolution of life science ontologies and mappings. In: International Workshop on Data Integration in the Life Sciences. pp. 11–27. Springer (2008)
8. Janik, M., Scherp, A., Staab, S.: The semantic web: collective intelligence on the web. Informatik-Spektrum **34**(5), 469 (2011)
9. Jiménez-Ruiz, E., Grau, B.C., Sattler, U., Schneider, T., Berlanga, R.: Safe and economic re-use of ontologies: A logic-based methodology and tool support. In: ESWC. pp. 185–199. Springer (2008)
10. Käfer, T., Abdelrahman, A., Umbrich, J., O'Byrne, P., Hogan, A.: Observing linked data dynamics. In: ESWC. pp. 213–227. Springer (2013)
11. Kamdar, M.R., Tudorache, T., Musen, M.A.: A systematic analysis of term reuse and term overlap across biomedical ontologies. Semantic web **8**(6), 853–871 (2017)
12. Meusel, R., Bizer, C., Paulheim, H.: A web-scale study of the adoption and evolution of the schema.org vocabulary over time. In: International Conference on Web Intelligence, Mining and Semantics. p. 15. ACM (2015)
13. Nishioka, C., Scherp, A.: Temporal patterns and periodicity of entity dynamics in the Linked Open Data Cloud. In: K-CAP. p. 22. ACM (2015)
14. Noura, M., Gyrard, A., Heil, S., Gaedke, M.: Concept extraction from the web of things knowledge bases. In: The International Conference WWW/Internet (2018)
15. Palma, R., Zablith, F., Haase, P., Corcho, O.: Ontology evolution. In: Ontology Engineering in a Networked World, pp. 235–255. Springer (2012)
16. Papakonstantinou, V., Fundulaki, I., Flouris, G.: Assessing linked data versioning systems: The semantic publishing versioning benchmark. International Workshop on Scalable Semantic Web Knowledge Base Systems@ISWC pp. 219 – 234 (2018)
17. Schaible, J., Gottron, T., Scherp, A.: Termpicker: enabling the reuse of vocabulary terms by exploiting data from the Linked Open Data Cloud. In: ESWC. pp. 101–117. Springer (2016)
18. Vandenbussche, P.Y., Atemezing, G.A., Poveda-Villalón, M., Vatant, B.: Linked Open Vocabularies (LOV): a gateway to reusable semantic vocabularies on the web. Semantic Web **8**(3), 437–452 (2017)
19. Zaki, M.J., Meira Jr, W., Meira, W.: Data mining and analysis: fundamental concepts and algorithms. Cambridge University Press (2014)