Towards Understanding the Evolution of Vocabulary Terms in Knowledge Graphs

Mohammad Abdel-Qader
Kiel University and ZBW – Leibniz Information Centre for Economics
stu120798@informatik.uni-kiel.de

Ansgar Scherp
ZBW – Leibniz Information Centre for Economics and Kiel University
a.scherp@zbw.eu

ABSTRACT

Vocabularies are used for modeling data in Knowledge Graphs (KG) like the Linked Open Data Cloud and Wikidata. During their lifetime, the vocabularies of the KGs are subject to changes. New terms are coined, while existing terms are modified or declared as deprecated. We first quantify the amount and frequency of changes in vocabularies. Subsequently, we investigate to which extend and when the changes are adopted in the evolution of the KGs. We conduct our experiments on three large-scale KGs for which time-stamped snapshots are available, namely the Billion Triples Challenge datasets, Dynamic Linked Data Observatory dataset, and Wikidata. Our results show that the change frequency of terms is rather low, but can have high impact when adopted on a large amount of distributed graph data on the web. Furthermore, not all coined terms are used and most of the deprecated terms are still used by data publishers. There are variations in the adoption time of terms coming from different vocabularies ranging from very fast (few days) to very slow (few years). Surprisingly, there are also adoptions we could observe even before the vocabulary changes are published. Understanding this adoption is important, since otherwise it may lead to wrong assumptions about the modeling status of data published on the web and may result in difficulties when querying the data from distributed sources.

1 INTRODUCTION

Vocabulary terms define the schema of Knowledge Graphs (KGs) such as the Linked Open Data (LOD) cloud or Wikidata. After ontology engineers built and published the first version of a vocabulary, the terms are subject to changes to reflect new requirements or shifts in the domains the vocabularies model. So far it is unknown how such vocabulary changes are reflected by the KGs that are using the terms. For example, data providers may not be aware that changes on the vocabulary terms happened, since it occurs rather rarely [6]. Explicitly triggering data providers to update their model is also challenging due to the distributed nature of KGs such as the LOD cloud. However, in general data providers may be interested in being notified when certain vocabulary term changes happen. But until today, data providers lack proper tools and services to track whether and what kind of changes on vocabulary terms happened. Likewise, ontology engineers lack a tool that reflects the adoption status of data w.r.t. their vocabularies and changes on the vocabulary terms.

We present a methodology and its implementation in a service to improve the understanding of vocabulary changes and how they are adopted in evolving KGs. Formally, we understand a vocabulary V as set of terms T. A term T is either a class C or a property

P. A set of terms relates to a vocabulary as $T(V) = C(V) \bigcup P(V)$. Changes on a vocabulary V are changes on its terms, i. e., the classes and properties. Data that uses classes and properties of a changed vocabulary should be updated according to those changes. This is needed especially, when a change is classified as deprecation. But there are also other types of changes like additions, modifications, or renaming a term. In a previous work [1], we manually conducted a qualitative analysis of vocabulary evolution on the LOD cloud. We analyze the changes for a set of vocabularies by clarifying which terms changed, the type of change, and if those changes were done on terms defined in the vocabularies or on the classes and properties that were imported from other vocabularies. Based on this qualitative analysis, we have developed a system to automatically extract changes on a set vocabularies and analyze the adoption time for the vocabulary changes, i. e., we monitor when changes on classes and properties are adopted by the KGs.

In this paper, we are studying the evolution of vocabulary terms in KGs. Our first research question is, when are the newly created terms of vocabularies adopted in KGs? Second, what is the usage rate of classes and properties for a set of vocabularies in each dataset, and are the Pay Level Domains which that highly use existing and deprecated terms? Third, are the deprecated terms are still used in KGs? We consider the two the basic types of changes, namely addition and deletion. Any other change like a modification can be expressed by the basic change types. As datasets, we use three well-known KGs. First, we use the Dynamic Linked Data Observatory (DyLDO) dataset [7], which are weekly snapshots for a set of Linked Data documents from the LOD cloud. Second, we analyze the Billion Triples Challenge (BTC)[1] datasets crawled from the LOD cloud in 2009, 2010, 2011, 2012, and 2014. From both datasets, we extract the Pay Level Domains (PLDs) that adopted changes of vocabulary terms. By PLD, we refer to any web domain that requires payment at a Top Level Domain (TLD) or country code TLD (cc-TLD) registrar [8]. As third dataset, we analyze the evolution of the Wikidata[2] KG. We extract the vocabulary terms from Wikidata and determine if changes on the vocabulary were done (additions or deprecations) and how these changes were adopted. Our experiments show that even if the frequency of vocabularies terms changes is rather low, they have a large impact on the real data. Furthermore, we found that most of the newly coined terms are adopted in less than one week after their publishing date. However, there are terms that are only adopted after several months or up to a few years after the date of creation. Moreover, there are also some terms adopted before their official publishing date. Additionally, many deprecated terms are still in use in KGs. Therefore we can claim that many

---

[1]http://challenge.semanticweb.org/, last accessed: June 29, 2017
[2]https://www.wikidata.org, last accessed: June 29, 2017

of those terms are in practice not really deprecated. Finally, we found that the percentages of the actually unused terms in KGs are more than 50% for most vocabularies, especially in the BTC dataset. For Wikidata, we extract which terms that are introduced to the vocabulary are actually never being used in the Knowledge Graph and aim to explain why there are terms that are not adopted. We found no deprecation of terms found in the Wikidata vocabulary. Moreover, there are 17 terms not used because they are for defining properties and their types. We think our work does help ontology engineers to select classes and properties that fit their needs by having a clear view about their adoption time and adoption rate, and help them in updating their own ontologies. In a subsequent step, we aim to provide this approach as freely available web service to track vocabulary changes and their adoption.

The remainder is structured as follows: In Section 2, we review the related work. Our methodology is presented in Section 3, followed by a description of the datasets and vocabularies extracted from these datasets in Section 4. Section 5 describes the results of our experiments. We discuss the findings and outcomes of our experiments in Section 6, before we conclude.

2 RELATED WORK

In terms of analyzing the use of structured data on the web, Meusel et al. [9] analyzed the evolution and adoption of schema.org. They made a comparison of the usage of schema.org terms over four years by extracting the structured data of the web pages that use this vocabulary from *WebDataCommon* Microdata datasets. They extracted only the quads whose object or predicate contains *schema.org*. They discovered that not all terms in schema.org are used and deprecated terms are still used, as it is also shown in this work. Furthermore, they found that publishing new types and properties is preferred over using schema.org's extension mechanism. Guha et al. [5] investigated the usage of the schema.org vocabulary in the structured data of a set of web pages. They analyzed a sample of 10 billion web pages crawled from *Google index* and *WebDataCommon*. They found that about 31% of those pages had some *schema.org* elements and estimated that around 12 million websites are using *schema.org* terms. For their analyses, the changes on vocabulary terms were not considered as it is investigated in this work. Furthermore, the authors said the wide usage of this vocabulary's terms was caused by the vast support from third-party tools, such as *Drupal* and *Wordpress*. Mihindukulasooriya et al. [10] conducted a quantitative analysis for studying the evolution of DBpedia, Schema.org, PROV-O, and FOAF ontologies. They proposed some recommendations such as the need of dividing large ontologies into modules to avoid duplicates when adding new terms. Furthermore, they suggest adding provenance information beside the generic metadata when the change occurred. Dividino et al. [4] analyzed how the usage of RDF classes and properties on the LOD cloud changed over time. Thus, they analyzed the combination of classes and properties that describe a resource but did not investigate whether a vocabulary and its terms have changed. The authors applied their analysis on a dataset of 53 weekly snapshots from the DyLDO dataset, as it is also investigated in this work. Furthermore, Dividino et al. [3] developed novel strategies for updating large-scale LOD datasets to keep local caches up-to-date. Two setups were

investigated. The first setup was a single-step update of a local data cache, i. e., taking only one snapshot as history into account. The second setup considered updates over a long period of time that used information about the total sum of data evolution as one of its features. The evaluation presented the performance for each setup based on bandwidth limitations. However, for both experiments, the changes on vocabulary terms were not considered as it is investigated in this work. Vandenbussche et al. [13] published a report that describes Linked Open Vocabularies (LOV). The report provides statistics about using LOV and its capabilities such as the total number of classes and properties and the top-10 search terms using LOV (from January 2015 to June 2015). But the report does not include information about adopting new terms and from which PLDs. Rathachari et al. [2] proposed a model that facilitates the understanding of organisms. Their model presents the changes in taxonomic knowledge in RDF form. The proposed model acts as a history tracking system for changing terms. But their model does not give information about how and when the terms are used, and which PLDs adopted changed terms. Schaible et al. [11] published a survey for the most preferred strategies for reusing vocabulary terms. The participants were 79 Linked Data experts and practitioners and were asked to rank several LOD modeling strategies. The survey concluded that terms that are widely used are considered as a better approach. Furthermore, vocabularies that are frequently used is a more important argument for reuse than the frequency of a single vocabulary term (ignoring the frequency of the vocabulary where the term belongs to). This survey can help to understand why there are some terms frequently used and why some of them are not used at all. Over six months, Käfer et al.[6] observed the documents retrieved from the DyLDO dataset they crawled. They analyzed those documents using different factors, their lifespan, the availability of those documents and their change rate. Also, they analyzed the RDF content that is frequently changed (triple added or removed). Finally, they observed how links between documents are evolved over time. Their study is important for the communities in different areas such as smart caching, link maintenance, and versioning. But their analyses does not include information about adopting new and deprecated terms.

3 ANALYSIS METHOD

Our analysis method consists of two steps: First, we determine vocabularies that have more than one published version on the web. Second, we investigate how the changed terms of vocabularies are adopted and used in the evolving KGs. For the first step, we based on Schmachtenberg et al. [12] who published a report with detailed statistics about a large-scale snapshot the LOD cloud. The snapshot comprises seed URIs from the datahub.io dataset[3], the BTC 2012 dataset[4], and the public-lod@w3.org mailing list[5]. We select a set of vocabularies that satisfy the following set of conditions and characteristics: (1) The vocabulary should have at least two versions published on the web to make a comparison between them. (2) We check if these two versions are covered by the dataset we investigate. For example, for the DyLDO dataset there is to be one

version of the vocabularies that has been published after May 6th, 2012. This is needed, since at this date the first snapshot of the DyLDO dataset has been crawled. (3) The vocabulary terms are to be directly used for modeling some data, i. e., there needs to be a direct use of the vocabulary terms by at least one triple in the published dataset. In contrast, vocabularies could also be just linked from a data publisher, where changes of external vocabularies may not have any impact on the published data.

Based on the criteria above, we examined 134 of the most used vocabularies listed in the state of the LOD cloud 2014 report by Schmachtenberg et al. [12]. We found 18 vocabularies that have more than one version. From them, 13 vocabularies have changes (additions or deprecations) on terms created by the ontology engineers of those vocabularies in the timeframe of the considered datasets. We downloaded the different versions of the extracted vocabularies using the Linked Open Vocabularies (LOV) observatory[6]. LOV is using a script that automatically checks for vocabulary changes in a daily basis and stores the detected version locally. We extracted the changes between each two successive versions of a vocabulary by using Protégé 4.3.0, which is a free, open source ontology editor application that allows users to create, edit, and compare ontologies[7].

Subsequently, we investigate in the second step how changed vocabulary terms are used in the evolving KGs. We extract all PLDs from the crawled triples that use the terms from the 13 vocabularies above. We use *Guava*[8] library version 16.0.1 to extract the PLD from URLs. We record, besides the date of the first appearance of a vocabulary term, also the number of triples that contain the term. This information is then used to compute the adoption time of vocabulary term changes over the dataset snapshots.

4 DATASETS AND VOCABULARIES

4.1 Datasets

We apply our analysis approach on three large-scale KGs. The first two are DyLDO and BTC and are obtained from the Linked Open Data cloud, and the third is Wikidata. Below, we briefly characterize the datasets.

The Dynamic Linked Data Observatory (DyLDO) is a repository to store weekly snapshots from a subset of web data documents [7]. They started crawling snapshots since May 2012. For our study, we parse 242 snapshots (from May 2012 until March 2017). The Billion Triple Challenge (BTC) is part of the Semantic Web Challenge[9], with a main goal to crawl a dataset from the LOD cloud. We used all available BTC datasets which were crawled in the years 2009, 2010, 2011, 2012 and 2014 to analyze the adoption of the extracted vocabularies in our study.

Wikidata is a knowledge base to collaboratively store and edit structured data. It is a free and multilingual repository[10]. In order to analyze the Wikidata vocabulary, we first extract the terms introduced by this vocabulary. Using the RDF Exports from Wikidata

page[11], we parse the terms and properties from the RDF dump files that were generated using the Wikidata toolkit[12]. We assume that the first snapshot of those files is the first version of the Wikidata vocabulary, and based on this assumption we parse the next dump files to extract the changes to the first version, and so on. Overall, there are 25 RDF dump files (from April 2014 until August 2016). Using those files, we extract the terms that are added or deprecated. Subsequently, we parse the Wikidata statements RDF dump files to extract the adoption of terms to analyze the adoption behavior for the Wikidata vocabulary's terms.

4.2 Extracted Vocabularies

The vocabularies that fulfill the conditions for the automated analysis of vocabulary terms evolution are listed in Table 1. The table also shows the number of versions of each vocabulary that are used in this study. Furthermore, it shows the total number of changes (additions and deletions) occurred during the vocabularies evolution.

Table 1: Overview of the vocabularies and their number of changes. The # changes column represents the total number of additions and deletions counted during the vocabulary evolution.

Vocabulary	# versions	# changes
Asset Description Metadata Schema (ADMS)[13]	2	18
Citation Typing Ontology (CiTO)[14]	3	218
The data cube vocabulary (Cube)[15]	2	6
Data Catalog Vocabulary (DCAT)[16]	2	13
A vocabulary for jobs (emp)[17]	2	1
Ontology for geometry (geom)[18]	2	2
The Geonames ontology (GN)[19]	7	31
The music ontology (mo)[20]	2	46
Open Annotation Data Model (oa)[21]	2	31
Core organization ontology (org)[22]	2	8
W3C PROVenance Interchange (Prov)[23]	5	168
Vocabulary of a Friend (voaf)[24]	4	8
An extension of SKOS for representation of nomenclatures (xkos)[25]	2	1

5 RESULTS

In Section 5, we summarize our findings after parsing the DyLDO and BTC datasets and extracting the required information for conducting this study. Section 5.1 lists our findings on changes of the 13 vocabularies in terms of additions and deprecations. Furthermore, we summarizes the results of usage and adoption of the 13 vocabularies in KGs datasets. Section 5.2 summarizes the findings that are related to the Wikidata vocabulary in terms of changes, usage, and adoption.

---

[6]http://lov.okfn.org/dataset/lov, last accessed: June 20, 2017

[7]http://protege.stanford.edu, last accessed: June 29, 2017

[8]https://github.com/google/guava/, last accessed: June 29, 2017

[9]http://challenge.semanticweb.org/, last accessed: July 10, 2017

[10]https://www.wikidata.org/, last accessed: June 29, 2017

[11]http://tools.wmflabs.org/wikidata-exports/rdf/exports.html, last accessed: June 29, 2017

[12]https://github.com/Wikidata/Wikidata-Toolkit, last accessed: June 29, 2017

5.1 Changes, Usage, and Adoption Time of Terms on the LOD Cloud

Figure 1 shows the 13 vocabularies in our study and the total number of classes and properties for each version of those vocabularies. Please note that most of them have an increased number of terms except two vocabularies. The number of classes and properties is decreased for the *ADMS* vocabulary, and *CiTO* had a huge decline in the number of classes.

During our analysis, we noticed that some of the deprecated properties are recreated again after they were deprecated. These recreated terms related to the *CiTO* and *GN* vocabularies. The *CiTO* vocabulary deprecated 18 properties in May 2014 (they were introduced in March 2010), and then they were recreated again in the version that was published in March 2015. The *GN* vocabulary was recreated three deprecated properties: *alternateName* (creation: October 2006, deprecation: September 2010, recreation: February 2012), *name* (creation: October 2006, deprecation: September 2010, recreation: October 2010), and *shortName* (creation: September 2010, deprecation: May 2010, recreation: February 2012). We can notice that *CiTO* has recreated those properties after around ten months of deprecation. Furthermore, for the *GN* vocabulary, two out of three deprecated terms are recreated in the new version after about 17 months, and the remaining one was recreated fast (after 13 days).

For each dataset, Table 2 shows the PLDs (based on all PLDs extracted during our analysis) that used terms related to the 13 vocabularies. For most vocabularies, we notice that there are no variety in the PLDs that use terms from the 13 vocabularies, i. e., there is one PLD that has the highest usage rate of terms. Furthermore, we can observe that *org* and *Cube* vocabularies are used by many PLDs. Table 3 shows an analysis for adopting the newly created terms for each vocabulary in our study. Additionally, it shows the total number of instances that the newly created terms appeared in. Furthermore, the table shows the mean (μ) and the standard deviation (σ) values in terms of days for adopting the newly created terms. We can notice that most of the newly coined terms are adopted, in less than 10 days. Furthermore, we notice that all new terms related to the *oa* and *xkos* vocabularies are not adopted at all. Table 4 shows the number of triples use terms from the 13 vocabularies. We can observe that *geom* does not appear in all the BTC datasets. Additionally, we can say that *emp* and *oa* vocabularies almost did not appear in the BTC and DyLDO datasets. Table 5 shows the PLDs with a high usage rate terms from the 13 vocabularies. We can notice that *geonames.org* has a high usage rate in the BTC2009 and BTC2010 datasets, while *dbtune.org* has a high usage rate in the BTC2014 and DyLDO datasets.

We noticed that most of the deprecated terms are still used after they were marked as deprecated. We found that *geonames.org* is the PLD which have the highest adoption rate for deprecated terms. The number of deprecated terms in all the 13 vocabularies and used by the *geonames.org* PLD, and extracted during our analysis of DyLDO, BTC2011, BTC2012, and BTC2014 datasets are 6, 6, 3, and 49, respectively.

Table 6 shows the percentage of unused terms from vocabularies depending on the extracted triples from the BTC and DyLDO datasets. We can notice that all terms of the *geom* vocabulary did not appear in BTC, beside most of the terms of *oa, emp*, and *CiTO*.

Table 2: A list of PLDs with the highest usage rate (based on all PLDs extracted during our analysis) that highly used terms related to the 13 vocabularies with changes. The (Amount) represent the total numbers of triples that the corresponding PLD appears in.

Vocabulary	BTC (Amount)	DyLDO (Amount)
ADMS	w3.org (5K)	w3.org (253K)
CiTO	ontologycentral.com (3.3M)	ontologycentral.com (4.7M)
Cube	ontologycentral.com (1.2M)	ontologycentral.com (1.8M)
DCAT	ontologycentral.com (5M)	ontologycentral.com (4.8M)
emp	purl.org (11K)	purl.org (14K)
geom	-	ign.fr (5.5K)
GN	geonames.org (90M)	geonames.org (40M)
mo	dbtune.org (84M)	dbtune.org (76M)
oa	w3.org (7K)	w3.org (13K)
org	data.gov.uk (21K)	w3.org (111K)
Prov	w3.org (62K)	dbpedia.org (1M)
voaf	purl.org (381K)	purl.org (9K)
xkos	270a.info (609K)	270a.info(4K)

Table 3: The adoption of newly created terms for each of the vocabularies in this study. The % of used terms column represent the percentage of the actually used terms compared to the total number of new created terms, and the # instances column represents the total amount of triples the newly created terms appeared in.

Vocabulary	% of used terms	# instances	μ (days)	σ (days)
ADMS	100%	31K	7	0
CiTO	100%	281K	7	0
Cube	100%	15K	7	0
DCAT	100%	104K	8.4	3.13
emp	100%	4K	7	-
geom	100%	16K	420	0
GN	100%	160M	127.76	255.33
mo	100%	45M	8.75	9.68
oa	0%	-	-	-
org	100%	173K	7	0
Prov	85%	121M	30.15	37.49
voaf	90%	75K	43.33	68.58
xkos	0%	-	-	-

In DyLDO, *CiTO* has the highest percentage of unused terms. Based on the triples extracted from DyLDO dataset, all the 37 created terms by *Cube* vocabulary appeared in the extracted triples. *org* vocabulary created 44 terms, and only 5 of them not appeared in the crawled triples of BTC dataset.
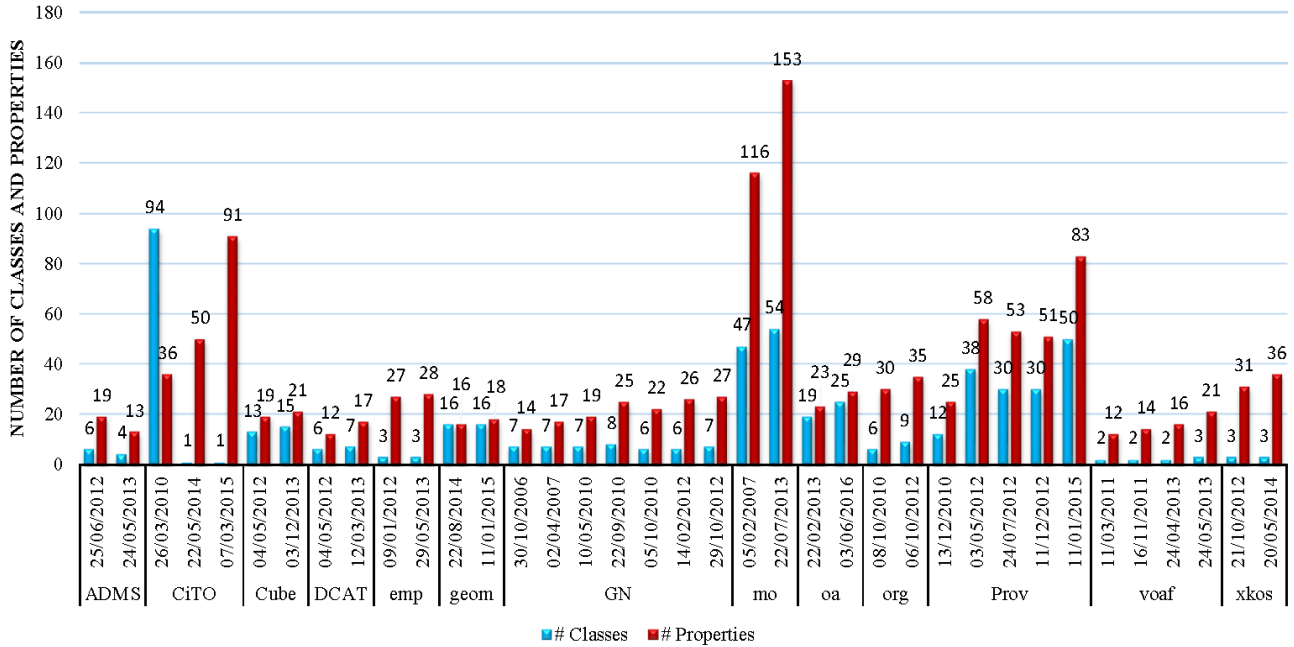
Figure 1: The total number of classes and properties for each of the 13 vocabularies versions in our study. The blue bar (left) represents the number of classes and the red bar (right) represents the number of properties.

Table 4: Number of triples in the DyLDO and BTC datasets that use terms from the 13 vocabularies.

Vocabulary	BTC09	BTC10	BTC11	BTC12	BTC14	DyLDO
ADMS	12	2	4	26	7K	337K
CiTO	0	0	4	4.5K	305K	1M
Cube	0	0	40K	8.4K	56M	12M
DCAT	0	0	12	4.5K	317K	9.8M
emp	0	0	0	0	238	26K
geom	0	0	0	0	0	5.7K
GN	81M	7.4M	477K	441K	1M	55M
mo	2M	1.7M	12M	4M	102M	83M
oa	0	0	0	0	192	23K
org	0	9	129	11K	20K	700K
Prov	0	9	43	902	3M	4M
voaf	0	0	2	0	4K	1M
xkos	0	0	0	0	610K	194K

Table 5: List of PLDs which have the highest usage rate (based on all PLDs extracted during our analysis) that used terms from the 13 vocabularies. The Repetition column represents the number of triples that the PLDs appears in.

Dataset	PLD	Repetition
BTC2009	geonames.org	81M
BTC2010	geonames.org	7M
BTC2011	zitgist.com	2.6M
BTC2012	rdfize.com	3.8M
BTC2014	dbtune.org	81.5M
DyLDO	dbtune.org	160M

5.2 Changes, Usage, and Adoption Time of Terms in Wikidata

After parsing the terms and properties from the RDF dump files for the period from April 2014 until August 2016, we extract the added and deprecated terms. Figure 2 shows the total number of classes and properties in each Wikidata snapshot. One can see, it grows to reach 11 classes and 27 properties in August 2017. Another important note is that there are no terms that are deprecated during the ontology evolution.

For the Wikidata vocabulary, ontology engineers added 3 classes and 9 properties during the analyzed period. The new classes are:

DeprecatedRank, PreferredRank, and *NormalRank.* And the properties are: *propertyTypeMonolingualText, propertyTypeProperty, rank, propertyQualifierLinkage, propertyReferenceLinkage, propertySimpleClaim, propertyStatementLinkage, propertyValueLinkage,* and *quantityUnit.*

Figure 3 shows the usage amount (in terms of triples) for newly created classes and properties on Wikidata. One can notice that the figure contains only 5 terms, while the actual number of newly coined terms is 12. This is because of there are 7 newly created terms are never used in Wikidata statements. Furthermore, we can notice that there is one class (*NormalRank*) and one property (*rank*) that have the highest usage amount of triples. Another interesting point is that the actually used terms from the newly created ones is adopted directly after their creation date. This note also can be applied to the other terms of the Wikidata vocabulary.

Table 6: Number and percentage of unused terms per vocabulary for the BTC and DyLDO datasets. The #terms column represents the total number of terms in the vocabulary.

Vocabulary	# terms	BTC	DyLDO	BTC%	DyLDO%
ADMS	31	21	1	68%	3%
CiTO	220	158	132	72%	60%
Cube	37	13	0	35%	0%
DCAT	23	11	2	48%	9%
emp	31	27	2	87%	6%
geom	34	34	1	100%	3%
GN	43	11	4	26%	9%
mo	208	74	4	36%	2%
oa	63	52	22	83%	35%
org	44	9	5	20%	11%
Prov	143	32	34	22%	24%
voaf	24	8	2	33%	8%
xkos	35	22	5	63%	14%
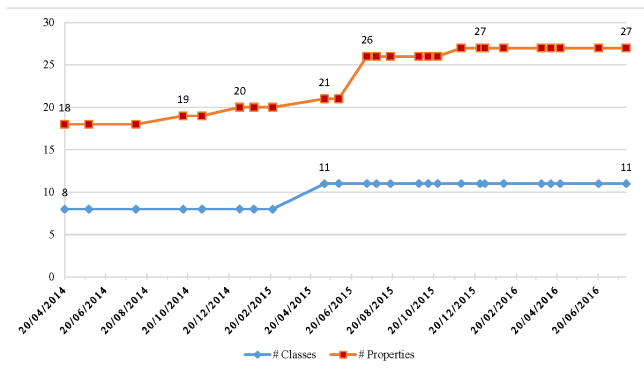
Figure 2: The total number of classes and properties in each RDF dump file generated by the Wikidata toolkit.
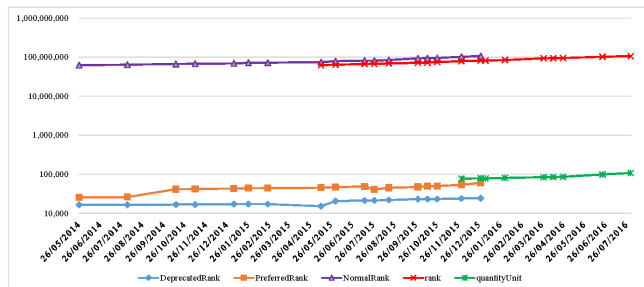
Figure 3: The usage amount in terms of triples for the newly created classes and properties on Wikidata vocabulary after parsing Wikidata RDF dump files.

6 DISCUSSION

In Section 6.1 we discuss the changes of terms of the 13 vocabularies in LOD datasets. Furthermore, we discuss the usage and adoption of the newly created terms. The Wikidata terms changes and adoption is discussed in Section 6.2.

6.1 Changes, Usage, and Adoption Time of Terms on the LOD Cloud

The number of additions and deprecations of terms are small. This is in line with existing studies[1, 5, 9]. There is one exception, the *CiTO* vocabulary's first version, which is published on March 2010, consists of 94 classes and 36 properties. After around four years of being online, the ontology engineers published an updated version that consisted of only one class and 50 properties. They removed all the 94 classes and replaced them with the new class *CitationAct*. Furthermore, most of the 36 properties of the first version are removed and replaced with new ones. The latest version of *CiTO* was published in March 2015. The number of classes did not change, but the number of properties increased to 91, and 18 of them were deprecated from the first version. We can say that the ontology engineers almost built a new ontology compared with the first published version, because they removed most of the terms and replaced them with new ones. This is important to note, since *CiTO* has grown much in popularity (e. g., the BTC 2014 dataset contained over 300K triples compared to 40K triples in BTC 2011). Another observation is that the *ADMS* vocabulary is the only vocabulary that has a decline in their terms (for both classes and properties).

For most vocabularies, we found that only one PLD has the highest usage rate of their terms, and there is no variety on the PLDs that use and adopt the vocabulary terms. In *ADMS*, *w3.org* was the PLD that highly used terms belonging to *ADMS*, but later (in 2015 and 2016) *deri.de* used the most of *ADMS* terms. On the other hand, there are some vocabularies that have been used from different PLDs. For example, *Cube* has been highly used from *ontologyCenter.com, esd.org.uk, linked-statistics.org*, and *linkedu.en.*

Analyzing the DyLDO dataset, the amount of usage remains in the same range of triples in most of the vocabularies except for the *mo* vocabulary (red line) and *Cube* vocabulary (light green line) as shown in Figure 4, which shows the amount of triples for each vocabulary. The triples are calculated on a weekly basis. In our opinion, these variations are depending on the aim of vocabularies, the generality, and the clearance of terms. *geonames.org* and *dbtune.org* are the PLDs (based on all PLDs extracted during our analysis) that use terms from all the 13 vocabularies. In the BTC2009 and BTC2010 datasets, *geonames.org* was the PLD that has a high usage rate of vocabulary terms. This is caused by the high number of triples using the *GN* vocabulary in those years. Furthermore, *dbtune.org* was the PLD with high usage rate in the BTC2014 and DyLDO snapshots from 2012 to 2014.

Although *geonames.org* was the PLD which have the highest usage rate for the 13 vocabularies, it is also the PLD that has the most uses deprecated terms. For example, in the BTC2011 dataset, *geonames.org* used six deprecated terms in about 522K triples. That number declined to three terms and about 181K triples in the BTC2012 dataset. But suddenly, in the BTC2014 the number of used deprecated terms has increased to be 49 terms, but in just about 5.5K triples. We assume that this was caused by not updating the documents and ontologies that used those deprecated terms. Figure 5 shows the usage of the term *gn:Country* in the DyLDO dataset which deprecated in September 2010. The figure shows that it had an increased usage rate until August 2015, then it declined and increased again to reach the peak point of usage in August 2016.
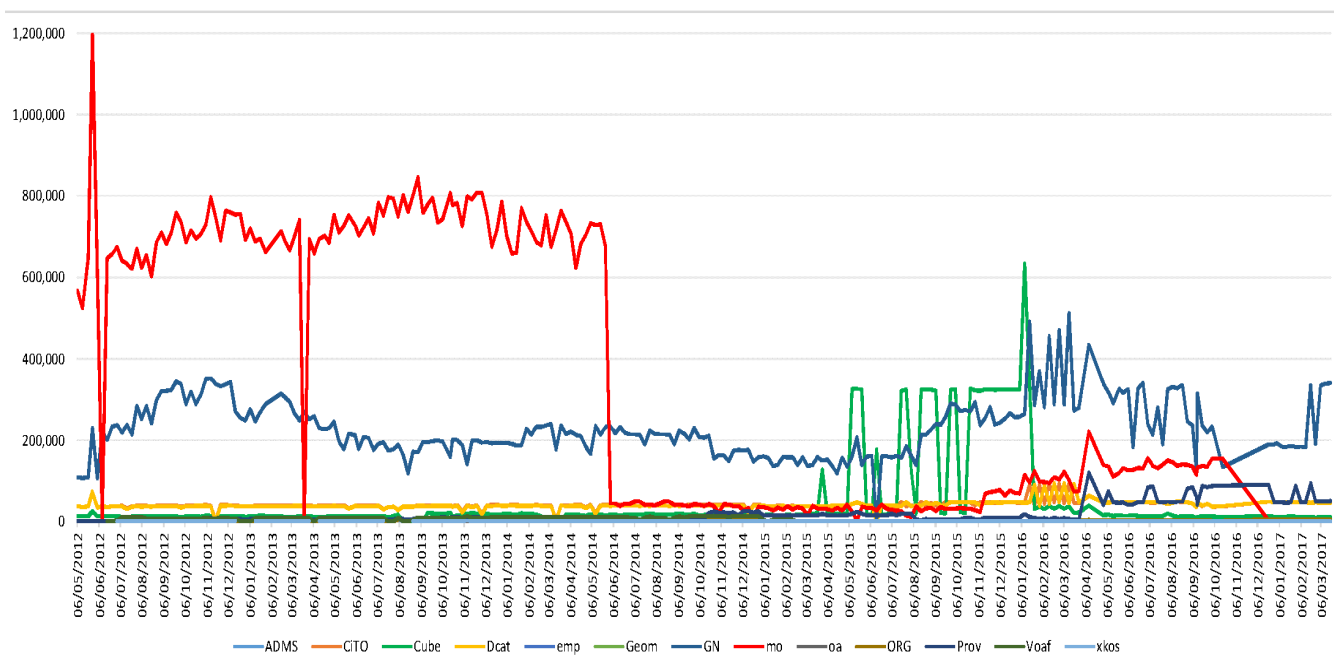
Figure 4: Amount of triples that had terms from the 13 vocabularies, in the DyLDO datasets. The X-axis represents the dates of DyLDO snapshots and Y-axis represents the total number of triples that use the terms of certain vocabularies.
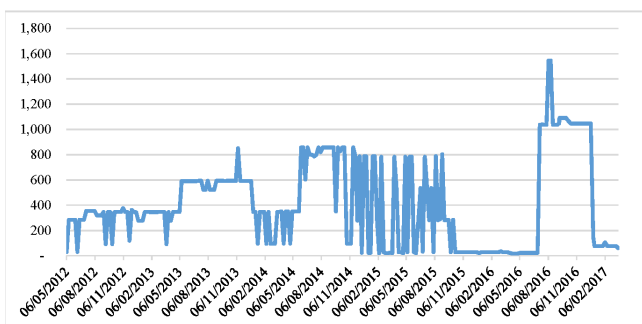

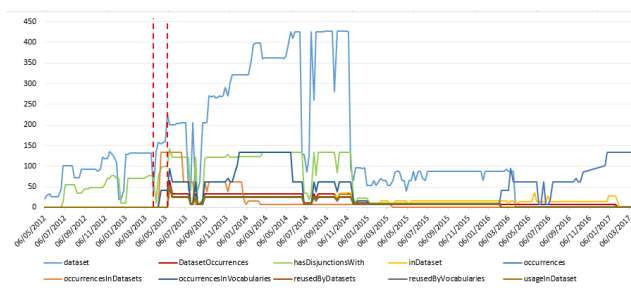**Figure 5: The usage of the class *gn:Country* in the DyLDO dataset after the class deprecated in September 2010.**

Figure 6: Amount of triples that use the *voaf* vocabulary in the DyLDO datasets. The vertical red dashed lines represent the time of publishing new versions for *voaf*.

Another interesting point is that some newly created terms are never adopted. For example, ontology engineers published a new version from *oa* vocabulary in June 2016 that has 21 new classes and properties. None of those terms are adopted (at least until April 2017, which is the last DyLDO snapshot we parsed). While the first version of *oa* is published in February 2013 with 42 terms, and all of them (except one term) are adopted in less than 3 months. We suggest that ontology engineers revise those unused terms.

Most of the newly coined terms are adopted directly (in less than one week). It is interesting that we found some terms adopted before their official publishing date. We can conclude that some of the new versions of vocabularies are unofficially online and can be used before the official publishing announcement. We assume that

in some cases, it may take time to finish the official procedures to publish the new version of the vocabulary.

The adoption of newly created terms are vary between vocabularies. For example, the *voaf* vocabulary has published four versions in March 2011, November 2011, April 2013, and May 2013. By parsing the DyLDO dataset, Figure 6 shows the usage amount in terms of triples for the new created terms only (which starts from November 2011 version). The *occurrenceInVocabularies* property (yellow line) is created in April 2013, and appears in 42 triples in the same month of its creation. The amount of usage goes in an increasing rate to reach its peak point of usage in May 2014. We think adding or removing terms and the increased number of available ontologies, have an impact on using those vocabularies. For more information about all the 13 vocabularies in this study, please refer to the Appendix.
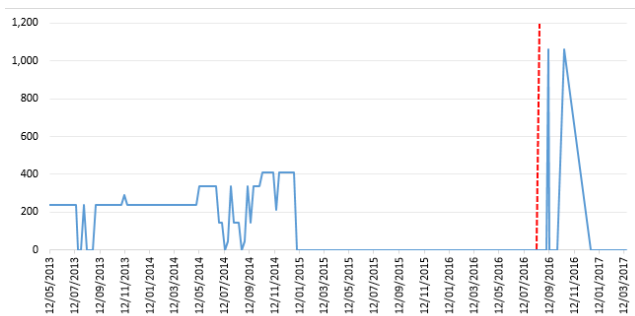
Figure 7: Amount of triples that use the *oa* vocabulary in the DyLDO datasets. The vertical red dashed lines represent the time of publishing new versions for *oa*.

Furthermore, the average number of days shows that most of the terms have fast adoption time, but in some vocabularies the mean value were more than 120 days, such as in *GN*. This average does not reflect the actual adoption behavior. From the new version of *GN* which newly created 21 terms, 17 terms are adopted within 7 days, and the remaining 4 terms are adopted in 600 and 650 days. Those four terms influenced in such result.

Publishing new versions of a vocabulary may influence the usage of terms. For example, Figure 7 shows the usage of *oa* vocabulary in the DyLDO dataset and when the new versions are published in June 2016. Even if the number of triples at its peak point (about 1K) maybe considered small, we can still notice an influence of these changes on the number of triples using the *oa* vocabulary. Thus, we assume that vocabulary changes such as adding or removing terms is a cause for updating data in KGs.

6.2 Changes, Usage, and Adoption Time of Terms in Wikidata

For the Wikidata vocabulary, an interesting point is that there are no deprecations of terms, even if they are not used during the time-frame investigated in this paper. Even if the terms are not used, ontology engineers does not mark them as deprecated (i. e., *Article* class). We found that the other unused terms are created to define properties and their types.

For the Wikidata vocabulary, like most of the LOD vocabularies, there is a small number of additions (3 classes and 9 properties), and no terms are deprecated. Beside that, there is a huge difference in the amount of triples that terms appeared in. For example, *NormalRank* and *Statement* classes have being used in about 106M and 81M triples, respectively. While the other classes (except the *Item* class) are used in less than 2.4M triples. The same observation can be made for properties, all of them appeared in less than 2.7M triples except the *rank* property, which introduced in May 2015 with appearance in about 62M triples, and in August 2016 the usage raised up to reach about 106M triples.

Another interesting point is that three classes (*DeprecatedRank*, *PreferredRank*, and *NormalRank*) suddenly disappeared from Wikidata statements after the snapshot in December 2015, and after about 8 months of usage (creation date was in May 2015). Another point is that there are two classes (*Article* and *Property*) and 15

properties which are not used in any statements of Wikidata. These are used to define the properties and their type (except the class *Article*).

7 CONCLUSION

Even small changes of vocabulary terms can have a deep impact on the real data that use those terms. For instance, *Prov* vocabulary remains in the range below than 25K triples before publishing its latest version in January 2015. After this version, the usage reached the level of 120K triples in 2016 and 100K in 2017. It is not surprising that deprecated terms are still in use. However, surprisingly some deprecated terms have been recreated after some time by their vocabularies. For instance, *CiTO* has recreated 19 terms after they were marked to be deprecated. Another example is *gn*, which recreated 3 deprecated terms in its latest versions.

REFERENCES

[1] Mohammad Abdel-Qader and Ansgar Scherp. 2016. Qualitative Analysis of Vocabulary Evolution on the Linked Open Data Cloud.. In *PROFILES@ ESWC*.

[2] Rathachai Chawuthai, Hideaki Takeda, Vilas Wuwongse, and Utsugi Jinbo. 2016. Presenting and preserving the change in taxonomic knowledge for linked data. *Semantic Web* 7, 6 (2016), 589–616.

[3] Renata Dividino, Thomas Gottron, and Ansgar Scherp. 2015. Strategies for efficiently keeping local linked open data caches up-to-date. In *International Semantic Web Conference*. Springer, 356–373.

[4] Renata Dividino, Ansgar Scherp, Gerd Gröner, and Thomas Grotton. 2013. Change-a-LOD: does the schema on the linked data cloud change or not?. In *Proceedings of the Fourth International Conference on Consuming Linked Data-Volume 1034*. CEUR-WS. org, 87–98.

[5] Ramanathan V Guha, Dan Brickley, and Steve Macbeth. 2016. Schema. org: Evolution of structured data on the web. *Commun. ACM* 59, 2 (2016), 44–51.

[6] Tobias Käfer, Ahmed Abdelrahman, Jürgen Umbrich, Patrick O'Byrne, and Aidan Hogan. 2013. Observing linked data dynamics. In *Extended Semantic Web Conference*. Springer, 213–227.

[7] Tobias Käfer, Jürgen Umbrich, Aidan Hogan, and Axel Polleres. 2012. Towards a dynamic linked data observatory. *LDOW at WWW* (2012).

[8] Hsin-Tsang Lee, Derek Leonard, Xiaoming Wang, and Dmitri Loguinov. 2009. IRLbot: scaling to 6 billion pages and beyond. *ACM Transactions on the Web (TWEB)* 3, 3 (2009), 8.

[9] Robert Meusel, Christian Bizer, and Heiko Paulheim. 2015. A web-scale study of the adoption and evolution of the schema. org vocabulary over time. In *Proceedings of the 5th International Conference on Web Intelligence, Mining and Semantics*. ACM, 15.

[10] Nandana Mihindukulasooriya, María Poveda-Villalón, Raúl García-Castro, and Asunción Gómez-Pérez. 2016. Collaborative Ontology Evolution and Data Quality-An Empirical Analysis. In *International Experiences and Directions Workshop on OWL*. Springer, 95–114.

[11] Johann Schaible, Thomas Gottron, and Ansgar Scherp. 2014. Survey on common strategies of vocabulary reuse in linked open data modeling. In *European Semantic Web Conference*. Springer, 457–472.

[12] Max Schmachtenberg, Christian Bizer, and Heiko Paulheim. 2014. Adoption of the linked data best practices in different topical domains. In *International Semantic Web Conference*. Springer, 245–260.

[13] Pierre-Yves Vandenbussche, Ghislain A Atemezing, María Poveda-Villalón, and Bernard Vatant. 2017. Linked Open Vocabularies (LOV): a gateway to reusable semantic vocabularies on the Web. *Semantic Web* 8, 3 (2017), 437–452.

A APPENDIX

A.1 Adopting the Newly Created Terms

In this section, we illustrate all the vocabularies with newly created terms and their adoption over time.
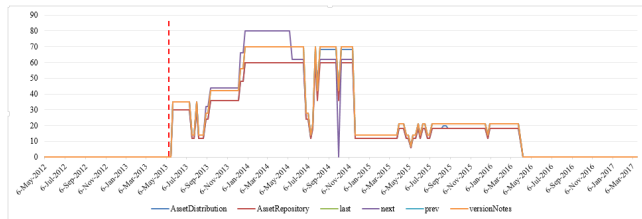

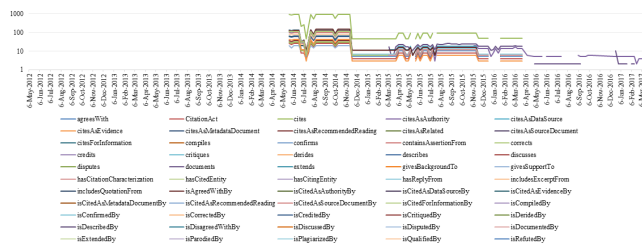**Figure 8: Adoption of terms from the Asset Description Metadata Schema (ADMS) vocabulary**

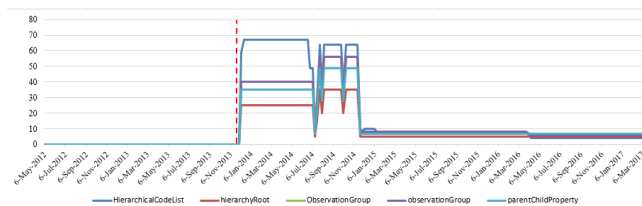**Figure 9: Adoption of terms from the Citation Typing Ontology (CiTO)**

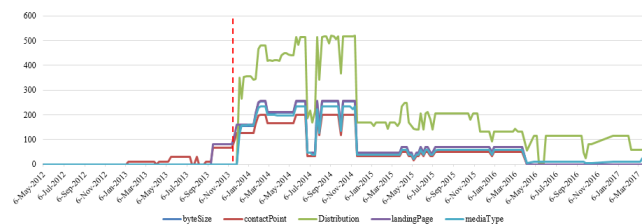**Figure 10: Adoption of terms from the data cube vocabulary (Cube)**

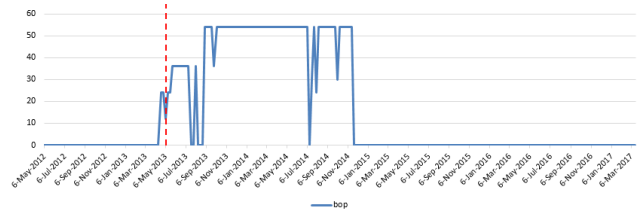**Figure 11: Adoption of terms from the Data Catalog Vocabulary (DCAT)**

Figure 12: Adoption of terms from the vocabulary for jobs (emp)
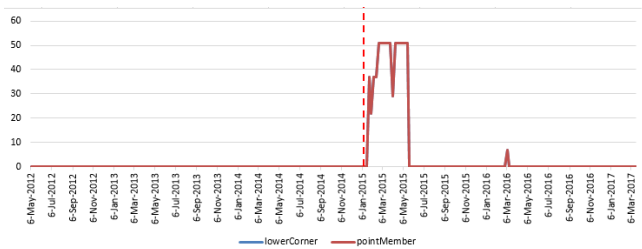
Figure 13: Adoption of terms from the ontology for geometry (geom)
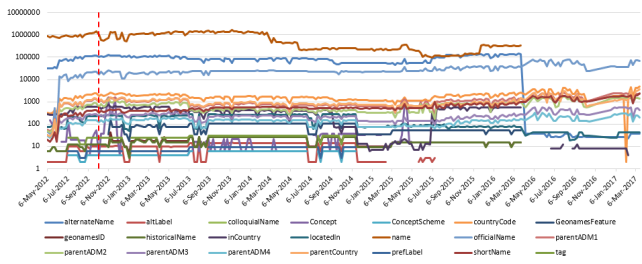
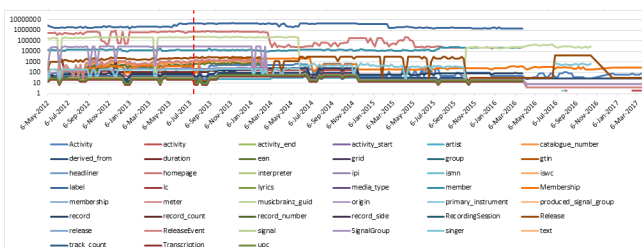**Figure 14: Adoption of terms from the Geonames ontology (GN)**

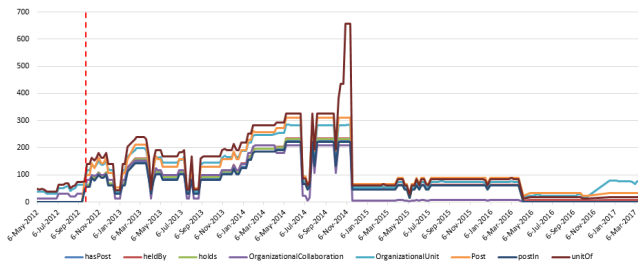**Figure 15: Adoption of terms from the music ontology (mo)**

Figure 16: Adoption of terms from the Core organization ontology (org)
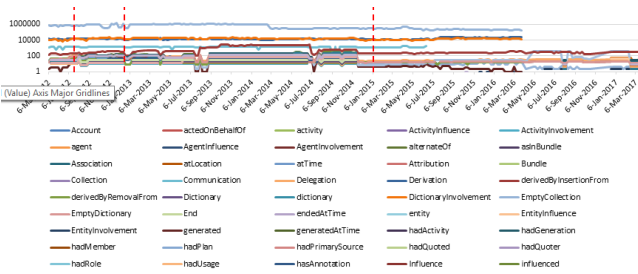

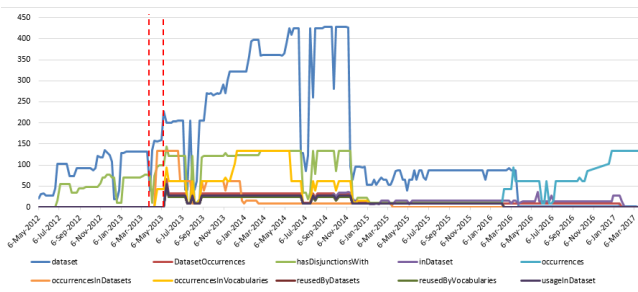**Figure 17: Adoption of terms from the W3C PROVenance Interchange (Prov) vocabulary**

Figure 18: Adoption of terms from the Vocabulary of a Friend (Voaf)