Supporting Web Vocabulary Development by Automated Quality Assessment: Results of a Case Study in a Teaching Context

Christian Mader[1] and Christian Wartena[2]

[1] University of Vienna, Faculty of Computer Science, Austria
christian.mader@univie.ac.at
[2] Hochschule Hannover, Germany
christian.wartena@hs-hannover.de

Abstract. Constructing controlled Web vocabularies such as thesauri for search and retrieval tasks is a widely intellectual process that relies on human experts. Thus, errors can occur which decrease the overall quality of the vocabulary. qSKOS is a tool that automatically checks Web vocabularies for potential quality problems and generates quality reports. In this paper we present the results of a case study designed to evaluate the impact of integrating qSKOS in the vocabulary creation process. It was carried out among students skilled in construction of controlled vocabularies. We collected in total 13 vocabularies in two versions and detected 15 different kinds of quality problems. For 11 of these problems we observed reduced occurrences after the vocabularies were revised by the participants based on the results of the generated quality report.

1 Introduction

Manual construction of controlled vocabularies (e.g., thesauri) is a tedious and time-consuming task that consists of multiple steps, such as defining subject area and scope, partitioning into subareas, collecting terms, imposing a structure on the terms and providing documentation. As structure and content evolves, the thesauri must furthermore be continually evaluated and revised.

When publishing thesauri on the Web using the Linked Data guidelines[1], they can serve as a means to connect datasets of different origins. The Simple Knowledge Organization System (SKOS[2]) is the de-facto standard for achieving this in a machine-friendly and human-readable way. It enables us to define a set of potential quality problems for controlled vocabularies that can be automatically evaluated and which we presented in our earlier work [9, 14]. We also found these quality problems to occur in existing vocabularies on the Web.

In courses in which thesaurus construction in general or more specifically SKOS is taught, it is hard to make students aware of problems that arise when working on large real-world thesauri. Often students make exercises only with

[1] http://www.w3.org/DesignIssues/LinkedData.html
[2] http://www.w3.org/2004/02/skos/

small vocabularies. Many of the typical problems arise when a thesaurus becomes larger and when it is not possible anymore to view the complete thesaurus structure at a glance. For the course teacher it will, on the other hand, become difficult and time consuming to identify all issues in a number of non-trivial thesauri.

Therefore, our hypothesis is that integration of automated quality checks in the thesaurus development process can (i) help students in a teaching environment to increase awareness for common quality problems and (ii) help developers to improve the quality of their thesauri.

In the following we report on the results of a case study performed in the spring term 2013 at the University of Applied Sciences in Hannover. Students with a background on thesaurus construction were required to develop thesauri covering businesses in different economic sectors. After a first development iteration, an automatically generated quality report was created and the results were handed out to the students. After a second development iteration, another quality report was generated. We compared the results of these reports and performed an in-depth analysis to find out about the quality issues that occurred and if and how they were addressed in the subsequent vocabulary version.

2 Background

Controlled Vocabularies are an effective way to organize knowledge by means of establishing semantic relationships between terms of the covered domain. Prominent examples of controlled vocabularies are, e.g., the Library of Congress Subject Headings[3] that are used to catalog materials held at the Library of Congress, or the Art & Architecture Thesaurus[4] for cataloging art collections developed by the Getty Research Institute. Recently, the principles of Linked Data have been widely recognized as an approach to publish controlled vocabularies on the Web in a machine-friendly, standardized format by using SKOS. It follows a concept-centric approach and is compatible with the ISO 25964 thesaurus standard[5].

2.1 Linked Data Validation

On the syntactical level, validators exist that help ensuring well-formedness of RDF data. More advanced tools[6] also take into account ontological information such as domain and range of OWL properties. In the field of ontology evaluation, approaches exist [11, 15, 16] which propose catalogs of patterns that can degrade an ontology's quality in terms of, e.g., understandability, validity or consistency. However, most of these patterns and validation checks cannot be applied to controlled vocabularies expressed using SKOS, because the SKOS schema imposes very few formal constraints.

[3] `http://id.loc.gov/authorities/subjects.html`
[4] `http://www.getty.edu/research/tools/vocabularies/aat/`
[5] `http://www.niso.org/schemas/iso25964/correspondencesSKOS/`
[6] e.g., DLLearner (`http://www.aksw.org/Projects/DLLearner.html`) or Pellet ICV (`http://clarkparsia.com/pellet/icv`)

2.2 Vocabulary Quality Assessment

In our earlier work we therefore proposed a set of semi-formally defined indicators for potential quality problems ("quality issues") specifically targeting SKOS vocabularies [9, 14]. We implemented them as computable quality functions in the qSKOS[7] tool and found frequent occurrences in vocabularies published on the Web. The quality issues are based on existing standards and guidelines for vocabulary development and linked data publication [1, 2, 6] as well as on a survey [8] conducted among experts in the field of controlled vocabulary development. However, despite positive reception and integration into a commercial product[8], we currently have no concrete numbers about how usage of the tool in a vocabulary development process influences occurrences of quality issues.

2.3 Approaches for Evaluating Quality Assurance Methods

Literature reporting on practical application and effectiveness of quality assurance methods in the creation process is still underrepresented. Coronado et al. [4] describes automated and manual quality assurance techniques applied on the editing and publication phase of the National Cancer Institute Thesaurus (NCIt). However, no figures on the actual number of found issues are provided. Gonalves et al. [5] provide a structural analysis of consecutive versions of the NCIt but do not focus on specific quality measures.

Concerning ontology evaluation, studies have been performed that use experiments to investigate feasibility and effectiveness of measures like, e.g., complexity or correctness. While we cannot directly apply these measures to controlled vocabularies, similar evaluation methodologies can be employed. Orme et al. [10] define ontology complexity and cohesion metrics like "Number of Properties" and "Average Fanout of Root Class" and perform an empirical analysis to evaluate them. Two experiments were performed to measure the correlation of the defined metrics with human perception of cohesion and complexity. In one experiment, 12 separate ontology instances were modified three times and reviewed by 18 evaluators. In our case study we follow a similar approach but let the vocabulary creators themselves decide about the feasibility of the quality report findings and incorporate changes respectively.

Strasunskas et al. [13] define measures for syntactic correctness and fitness of an ontology in a search task and evaluate their findings in an experiment among 21 students working with four different ontologies in two different versions. The students were divided into two groups and required to perform search tasks with subsequent judgment of the relevance of the results. Contrasting our approach, also in this setup the students had an observing role and no possibility to change the used ontologies.

---

[7] Available under an open-source license at `https://github.com/cmader/qSKOS/`

[8] `http://www.poolparty.biz/portfolio-item/poolparty-thesaurus-server/`

3 Materials and Methods

3.1 Quality Analysis

We used version 0.9.5 of qSKOS which checks for 21 potential quality issues. Checking for missing incoming links has been omitted in this study because the created vocabularies were not published online, leading to 20 checked issues. All submitted vocabularies and the generated reports can be retrieved online[9]. In total we analyzed 26 vocabularies, 13 for each submission containing between 43 and 111 concepts. Table 1 provides a description of all quality issues checked by qSKOS that were observed in this case study. An in-depth coverage of the quality issues can be found in our earlier work [9, 14].

Table 1. Description of observed quality issues.

Quality Issue Name	**Description**
Orphan Concepts	Concepts that are not semantically related to other concepts
Relation Clashes	Concepts being both hierarchically and associatively related
Overlapping Labels	Identical labels for distinct concepts
Inconsistent Preferred Labels	Concepts with more than one `skos:prefLabel` per language tag
Omitted or Invalid Language Tags	Missing language information for labels and textual descriptions
Mapping Clashes	Concepts related by both `skos:exactMatch` and one of `skos:broadMatch` or `skos:relatedMatch`
Disjoint Labels Violations	Identical preferred, alternative and hidden labels for one concept
Valueless Associative Relations	Associatively related concepts sharing the same broader concept
Disconnected Concept Clusters	Clusters of concepts that are not connected to each other
Cyclic Hierarchical Relations	Concepts related by mutual `skos:broader` relations
Broken Links	Links that do not resolve with an HTTP status 200 after redirections
Undocumented Concepts	Concepts lacking documentary properties such as `skos:note`
Incomplete Language Coverage	Concepts lacking labels in a language that is present for other concepts
Unidirectionally Related Concepts	No reciprocal relations, e.g. only one of `skos:broader` or `skos:narrower` is used
Missing Out-Links	Concepts not being linked to other resources on the Web

3.2 Data Acquisition and Studied Vocabularies

Construction of a SKOS vocabulary was an obligatory part of both the bachelor's and the master's course. For the bachelor students this was a classroom exercise, that was done in small groups of two or three students. The master students constructed a thesaurus as an individual exercise as part of the examination.

Students from the master course had to select an economic sector and find websites of 20 companies in this sector. All selected companies are located in

[9] `http://tinyurl.com/mv8vocs`

Germany and have a German web site. In the next step we collected characteristic words from these websites. These lists of words then should form the base for the construction of a thesaurus. Students had the freedom to remove irrelevant words from the list and add important missing terms. Table 2 lists all chosen domains alongside with their German reference that will be used throughout this paper.

Table 2. Thesaurus domains with German references.

Domain	**Reference**	**Domain**	**Reference**
Fashion and Clothing Industry	Bekleidung	Pharmacy	Pharma
Goat Farms	Ziegenhof	Library Information Systems	Bibliothekssoftware
Wind Energy	Wind	Mechanical Engineering	Maschinenbau
Shipbuilding	Werften	Orthopedic Technology	Orthopaedie
Paper and Board Industry	Papierindustrie	Medical Technology	Medtechnik

We used crawler4j[10] to crawl the websites. For a few companies the crawling was not successful and no pages could be retrieved. Since some companies have a site with a high number of pages, we limited the number of pages to be retrieved to 120. The limitation serves the practical goal of keeping the size of the corpus moderate, but also has more fundamental reasons: We expect that even a large company should be described rather well on the first two levels of a web site. If we crawl in a breadth first way, as we do, we might expect that at some point we have seen the core information of a company. If more pages follow, we might get more and more specific information on detailed topics, that even could obscure the more important and central information. The limit of 120 is rather arbitrary and turned out to be a size that allows us for almost all companies in our list to crawl the complete site. In total 14 673 pages were retrieved, which averages to 70.5 pages per company, with a total amount of about 4.6 million words. Almost the same corpus was used for the keyword extraction experiments described in [17].

We did not do any boiler-plate removal since it turned out that in many cases relevant and interesting words would be removed. E.g., a list of products or departments is often given as a menu, that might be removed. The whole corpus is tokenized and all words are lemmatized and tagged with their part of speech by the TreeTagger [12].

Each of the sub-corpora, i.e. the texts collected form the websites of the companies belonging to one economic sector, was used to construct a list of domain-specific terms. As candidates for the thesaurus, all words are selected that (i) are tagged as common noun, (ii) occur at least 5 times in the sub-corpus and (iii) have a relative frequency that is higher than the frequency in the general DeWaC [3] corpus.

[10] `http://code.google.com/p/crawler4j/`

Due to the fact that we did not use boiler-plate removal there are a lot of single words that are not part of a well formed sentence. Consequently a lot of errors are made by the part-of-speech tagger and a number of words that are not common nouns end up in the lists of term candidates. Moreover, many words are included, that are not typical for the economic sector of the sub-corpus, but for website texts. This is mainly a consequence of the completely different strategies for collecting texts for our corpus and for the DeWaC reference corpus. Despite these errors, each of the lists contained enough relevant words that could serve as a basis for a domain specific thesaurus.

3.3 Vocabulary Construction

After completion of the list of candidate terms each of the master students could work on the thesaurus on the chosen economic sector. Students were instructed to use all relevant words from the generated list, where judgment of relevance was left to the students own opinion. Moreover they were allowed to add a limited number of terms, if these terms are necessary to construct the thesaurus. Two criteria were mentioned explicitly: (i) a term might be necessary to represent an otherwise missing intermediate level or natural more general term to a number of more specific terms, and (ii) if otherwise a more general term would only have only one daughter. Finally, they were instructed to find matching terms in another thesaurus for at least 10 terms.

Since TopBraid Composer[11] was used throughout the course, all students used this tool for the thesaurus construction as well. In order to check the quality of the thesaurus they could send it once to the course teacher in order to get the qSKOS report. The differences between this first submission and the thesaurus that they finally submitted as a part of their assignment, is used below to get insight in the value of qSKOS.

The bachelor students were given the same instructions, but without the requirement to find matching terms in other vocabularies. These students worked in small groups of two or three students on two afternoons in a classroom setting. The lists of words used by these students were the same as those used by the master students. Only three thesauri of this course are included in this study. Some thesauri could not be used since the students have used qSKOS as a tool to continuously improve the thesaurus quality. Furthermore, some thesauri did not achieve a level of maturity that allows for a useful application of qSKOS.

4 Results

We counted the number of vocabularies that show a specific quality issue before qSKOS quality assessment (QA) and afterwards. Figure 1 shows that 11 of all 20 assessed quality issues lie in the right side of the dotted line, i.e., after the quality check less vocabularies were affected by these issues than before. Eight

[11] `http://www.topquadrant.com/products/TB_Composer.html`

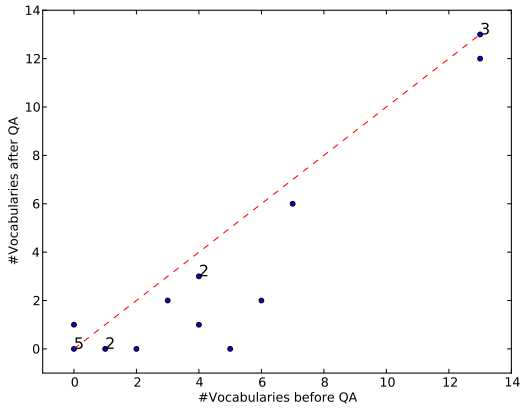

Fig. 1. Number of vocabularies affected by quality issues before and after QA.

issues either did not occur in any vocabulary or occurred in all vocabularies and did not improve after QA. One issue is positioned at the left side of the line, because it did not occur in any vocabulary before QA but showed up in one vocabulary afterwards. Table 3 shows issue occurrences in more detail.

To get a more detailed impression of the way the vocabularies were influenced by each quality issue, we calculated the issue occurrences for each vocabulary before and after QA. Figure 2 shows the number of vocabularies where less (improvements), more (degradation) or equal (no change) issue occurrences were spotted in the revised version.

In the following we elaborate on the quality issue changes in detail. Our findings can be summarized as follows:

- 13 of 20 checked quality issues were improved in at least one vocabulary.
- Seven quality checks led to improvements in up to five vocabularies.
- For eight other quality checks, besides improvements in up to seven vocabularies also degradations in up to four vocabularies were found.

4.1 Improvements without Degradations

Orphan Concepts were resolved in all five vocabularies where they occurred. The affected concepts were either removed from the vocabularies or hierarchically related to existing concepts. In some cases (Pharma) concepts where merged in a way that they became alternative labels of other concepts and were removed from the vocabulary afterwards.

Relation Clashes were resolved by the contributors for four vocabularies. In three vocabularies all occurrences of this issue were fixed and for one vocabulary (Medtechnik) the occurrences could be reduced from 25 to six. The applied

Table 3. Occurrences of quality issues before and after QA.

Quality Issue	Before QA	After QA	Difference
Orphan Concepts	5	0	5
Valueless Associative Relations	6	2	4
Relation Clashes	4	1	3
Inconsistent Preferred Labels	2	0	2
Omitted or Invalid Language Tags	3	2	1
Incomplete Language Coverage	7	6	1
Undocumented Concepts	13	12	1
Overlapping Labels	4	3	1
Disconnected Concept Clusters	4	3	1
Mapping Clashes	1	0	1
Disjoint Labels Violation	1	0	1
Solely Transitively Related Concepts	0	0	0
Omitted Top Concepts	0	0	0
Top Concepts Having Broader Concepts	0	0	0
Missing Out-Links	13	13	0
Broken Links	13	13	0
Undefined SKOS Resources	0	0	0
Unidirectionally Related Concepts	13	13	0
HTTP URI Scheme Violation	0	0	0
Cyclic Hierarchical Relations	0	1	-1

resolution strategies were to remove the associative relations (Maschinenbau, Wind), change them to reference other concepts (Wind2) or replace them with a hierarchical relation (Wind2). Some clashes (Medtechnik) were resolved by changing the hierarchical structure of the affected concepts. However, we assume that these substantial changes of the Medtechnik vocabulary led to introduction of the remaining 6 occurrences of this issue.

Occurrences of *Overlapping Labels* could be observed in four vocabularies and were improved in two of them. For one vocabulary (Bibliothekssoftware), all four occurrences were addressed by rephrasing preferred labels and removing alternative labels. The other improved vocabulary (Papierindustrie) initially showed conflicts between preferred and alternative labels of two concepts but only the latter were addressed in the subsequent version.

Inconsistent Preferred Labels occurred in two vocabularies and were fixed in both of them by either removing or rephrasing conflicting labels or conversion to alternative labels.

Three vocabularies were affected by *Omitted or Invalid Language Tags*, but two vocabularies showed no change and in only one vocabulary (Orthopaedie) this issue was improved. In this vocabulary, all issue occurrences were fixed by adding language tags to the three `rdfs:labels` where they were missing.

Mapping Clashes were observed in only one vocabulary (Wind2) with one concept which was mapped by both `skos:exactMatch` and `skos:broadMatch` to the same "external" resource. The issue was resolved by removing the latter. The same vocabulary was the only one which showed *Disjoint Labels Violations* for one concept, which were fixed by rephrasing the preferred label.

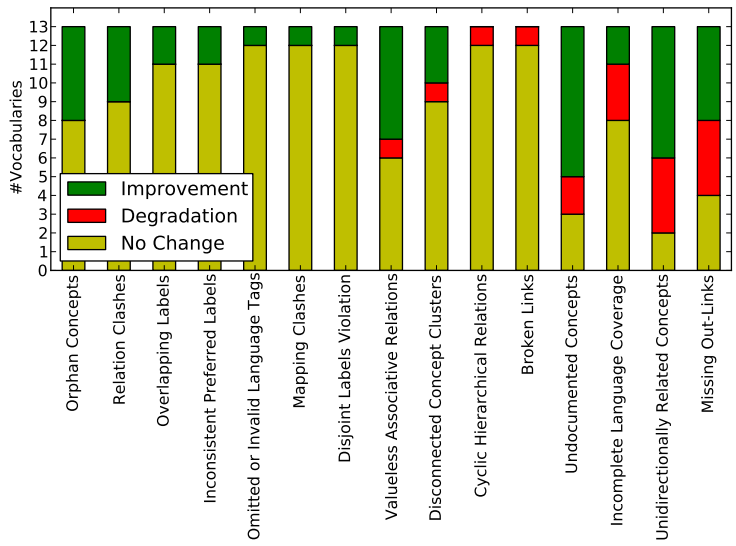

Fig. 2. Quality changes by issue.

4.2 Improvements with Degradations

Valueless Associative Relations were observed in six vocabularies before QA. In five vocabularies, all occurrences were fixed, in one vocabulary (Medtechnik) all but one occurrence was fixed. This unfixed issue did not occur in the initial version of the vocabulary, thus it has been introduced by changes in the hierarchical structure of the improved vocabulary. One vocabulary (Wind2) in the first version did not show any valueless associative relations, but one such relation was introduced by addition of a hierarchical relation.

Disconnected Concept Clusters were observed in four vocabularies before QA and improved in three of them. After the quality check the number of disconnected concept clusters were reduced to one "giant component" in two vocabularies. In one vocabulary (Pharma) the number of disconnected concept clusters decreased substantially from 18 to five. However, one vocabulary introduced an additional disconnected concept clusters after the check that was not present before. This cluster defines a hierarchical branch of materials and consists of eight of the 12 new concepts that have been newly introduced after QA.

None of the vocabularies showed *Cyclic Hierarchical Relations* before initial QA. However, in one vocabulary (Bekleidung) one cycle was identified in the final version. It seems that introduction of this issues was a side-effect when reducing the number of disconnected concept clusters from three to one. One concept in a cluster of only two concepts was hierarchically reorganized which caused the cycle.

In each of the created vocabularies we were able to spot at least one *Broken Link*. They were caused by the XML root namespace definition (set to, e.g., `http://hs-hannover.de/maschinenbau#`) for the created concepts. These links did not resolve because the vocabularies were not published online. One vocabulary (Papierindustrie) contained three external links to DBpedia and AGROVOC[12] which returned an HTTP status other than 200. After the check, another link to DBpedia was introduced that also did not resolve.

Undocumented Concepts occurred in all vocabularies and generally changed very few after QA. The changing numbers are mainly caused by concepts that were removed or newly added. However, in four vocabularies we actually noted intentional manual additions of `skos:scopeNote`s and `skos:definition`s. In one vocabulary (Maschinenbau2) all but four concepts were undocumented and these undocumented concepts were fixed after QA with `skos:definition`s.

Incomplete Language Coverage was spotted in seven vocabularies by QA and improved in two vocabularies in the final version. In one vocabulary (Maschinenbau2), two English labels were provided for concept that only had German labels. Another vocabulary (Werften) initially contained 80 concepts with German labels. Only two concepts were labeled in English. In the final version, the English alternative label was removed and the language tag of the preferred label "Cruises" was changed from `@en` to `@de`. These changes can actually be considered a degradation because correct information in the original vocabulary was removed and changed to be incorrect. In three other vocabularies, minor degradations could be observed that were caused by additions and removals of new concepts and labels as it was also the case for the issue of *Undocumented Concepts*.

Unidirectionally Related Concepts were explicitly not required to be treated by the experiment contributors. Thus, for all vocabularies that showed issues of this kind before the quality check, they remained nearly constant with changes only being indirectly caused by concept additions and removals.

Improvements for *Missing Out-Links* could be observed in five vocabularies. In three of them (Wind2, Werften, Pharma), the experiment contributors deliberately introduced mapping relations to external resources on the web. Removals and additions of concepts led to the side-effect of improvements of this issue in two vocabularies and degradations for another two vocabularies. For some reason in two vocabularies (Bekleidung2, Medtechnik) seemingly correct mapping relations were removed by the participants.

5 Discussion

Our study showed that quality issues were found in each of the vocabularies created by the participants. However, not all quality issues that we checked against were also observed. Only 15 of 20 kinds of quality issues did occur in the vocabularies because of two reasons: Either some SKOS constructs like concept

[12] `http://aims.fao.org/standards/agrovoc/about`

schemes or top concepts were not used by the participants or, in the case of *Undefined SKOS Resources* and *HTTP URI Scheme Violation*, the workflow of TopBraid Composer prevented such kind of errors.

From the 15 occurring kinds of quality issues, 13 of them were reduced in at least one vocabulary after the quality check. However, for eight kinds of quality issues we also noted increased occurrences (i.e. degradations) which were mainly caused by side effects of other changes like, e.g., addition or removal of concepts or changes in the hierarchical structure. These degradations could have probably been reduced by requiring the participants to perform a quality check before finally submitting the vocabulary.

The mentioned side-effects also show a weakness of our approach of counting the number of improvements for each quality issue: For *Unidirectionally Related Concepts*, we could observe a relatively high number of both improvements and degradations although participants were instructed not to take this issue into account. However, detailed examinations of the vocabularies and used issue resolution strategies show that the majority of all improvements were caused by the participants in explicitly resolving the respective quality issues.

In one case (Werften) changes were apparently introduced to achieve improvements for *Incomplete Language Coverage* although they actually cause a loss of information in favour of "better" issue occurrence values (i.e., a resulting value of zero for this quality issue). The reason for this kind of change could have been a misinterpretation of the experiment's goal by the participants and concerns of achieving a lower grade if some issues are not completely fixed.

6 Conclusion and Future Work

In this article we reported on the results of integrating automatic quality assessment into a controlled vocabulary development process. We found improvements for 13 of 15 kinds of quality issues after presenting a quality report to the participants. Due to the obtained results we conclude that automated quality assessment can be used to provide valuable feedback for identifying common mistakes that otherwise would not have been noted. Furthermore, we also observed that quality issues have been newly introduced in the second version of some thesauri that were not present before. To reduce these degradations, we consider it useful to run quality checks on a regular basis, providing continuous feedback to the developers. In addition, we plan to improve the readability of qSKOS reports for non-expert users.

We are aware that due to the limited number of participants further studies are required to increase confidence in the data. Therefore we plan to set up further experiments with a higher number of participants that also include experts in controlled vocabulary development. Furthermore, we plan to study the effects of continual automated quality checks by integrating them into tools like Databugger [7] or Semantic Web notification frameworks[13].

[13] like, e.g., Resource SubscrIption and Notification sErvice (rsine, `https://github.com/rsine/rsine`)

References

1. ISO 25964-1: Information and documentation – Thesauri and interoperability with other vocabularies – Part 1: Thesauri for information retrieval. Norm, International Organization for Standardization (2011)
2. Aitchison, J., Gilchrist, A., Bawden, D.: Thesaurus construction and use: a practical manual. Aslib IMI (2000)
3. Baroni, M., Bernardini, S., Ferraresi, A., Zanchetta, E.: The WaCky Wide Web: A Collection of Very Large Linguistically Processed Web-Crawled Corpora. Language Resources and Evaluation 43 (3): 209-226 43(3), 209–226 (2009)
4. de Coronado, S., Wright, L.W., Fragoso, G., Haber, M.W., Hahn-Dantona, E.A., Hartel, F.W., Quan, S.L., Safran, T., Thomas, N., Whiteman, L.: The nci thesaurus quality assurance life cycle. Journ. of Biomed. Informatics 42(3), 530–539 (2009)
5. Gonalves, R.S., Parsia, B., Sattler, U.: Analysing the evolution of the nci thesaurus. In: CBMS. pp. 1–6. IEEE (2011)
6. Hogan, A., Harth, A., Passant, A., Decker, S., Polleres, A.: Weaving the pedantic web (2010)
7. Kontokostas, D., Westphal, P., Auer, S., Hellmann, S., Lehmann, J., Cornelissen, R.: Databugger: A test-driven framework for debugging the web of data. In: Proc. of the 23rd international conference on World Wide Web (2014), to appear
8. Mader, C., Haslhofer, B.: Perception and relevance of quality issues in web vocabularies. In: I-SEMANTICS 2013. Graz, AUT (2013)
9. Mader, C., Haslhofer, B., Isaac, A.: Finding quality issues in skos vocabularies. In: TPDL 2012 Therory and Practice of Digital Libraries. Germany (May 2012)
10. Orme, A.M., Yao, H., Etzkorn, L.H.: Indicating ontology data quality, stability, and completeness throughout ontology evolution. Journal of Software Maintenance 19(1), 49–75 (2007)
11. Poveda-Villaln, M., del Carmen Surez-Figueroa, M., Gmez-Prez, A.: Validating ontologies with oops! In: ten Teije, A., Vlker, J., Handschuh, S., Stuckenschmidt, H., d'Aquin, M., Nikolov, A., Aussenac-Gilles, N., Hernandez, N. (eds.) EKAW. Lecture Notes in Computer Science, vol. 7603, pp. 267–281. Springer (2012)
12. Schmid, H.: Improvements in Part-of-Speech Tagging with an Application to German. In: Proceedings of the ACL SIGDAT-Workshop (1995)
13. Strasunskas, D., Tomassen, S.L.: Empirical insights on a value of ontology quality in ontology-driven web search. In: Meersman, R., Tari, Z. (eds.) OTM Conferences (2). Lecture Notes in Computer Science, vol. 5332, pp. 1319–1337. Springer (2008)
14. Suominen, O., Mader, C.: Assessing and improving the quality of skos vocabularies. Journal on Data Semantics (2013), accepted, in press
15. Tao, J., Ding, L., McGuinness, D.L.: Instance data evaluation for semantic web-based knowledge management systems. In: HICSS. pp. 1–10. IEEE Computer Society (2009)
16. Vrandecic, D.: Ontology Evaluation. Ph.D. thesis, KIT, Fakultät für Wirtschaftswissenschaften, Karlsruhe (2010)
17. Wartena, C., Garcia-Alsina, M.: Challenges and Potentials for Keyword Extraction from Company Websites for the Development of Regional Knowledge Maps. In: Liu, K. (ed.) Proceedings of the International Conference on Knowledge Discovery and Information Retrieval and the International Conference on Knowledge Management and Information Sharing, pp. 241–248. Scitepress (2013)