Towards Ontology Quality Assessment

Silvio Mc Gurk[1], Charlie Abela[1], and Jeremy Debattista[2]

[1] Department of Intelligent Systems, Faculty of ICT, University of Malta, Malta,
silvio.mcgurk.15@um.edu.mt, charlie.abela@um.edu.mt
[2] Enterprise Information Systems, Fraunhofer IAIS / University of Bonn, Germany,
debattis@cs.iai.uni-bonn.de

Abstract. The success of systems making use of ontology schemas depend mainly on the quality of their underlying ontologies. This has been acknowledged by researchers who responded by suggesting metrics to measure different aspects of quality. Tools have also been designed, but determining the set of quality metrics to use may not be a straightforward task. Research on ontology quality shows that detection of problems at an early stage of the ontology development cycle is necessary to reduce costs and maintenance at later stages, which is more difficult to achieve and requires more effort. Assessment using the right metrics is therefore crucial to identify key quality problems. This ensures that the data and instances of the ontology schema are sound and fit for purpose. Our contribution is a systematic survey on quality metrics applicable to ontologies in the Semantic Web, and preliminary investigation towards methods to visualise quality problems in ontologies.

Keywords: ontology quality metrics, ontology engineering, ontology evaluation, quality visualisation

1 Introduction

Many ontologies have been designed and developed over time, spanning a number of domains and including a number of concepts. Ontologies have been used in various domains including gene ontologies [2] and as unification tools in biomedicine [17], in education to enhance learning experiences [19] and in information retrieval systems [4]. As ontologies are being developed and reused, the need to address quality issues becomes an important factor as having a true understanding of the quality of an ontology helps future data publishers to choose ontologies based on 'fitness for use' [13]. Extensive research has been carried out along the years to help identify quality problems in ontologies [7, 23, 20, 3, 21, 10, 18, 11]. As a result of this research, a number of quality metrics have been suggested. These are coupled with tools and quality frameworks [5, 15, 7, 23, 25, 21] that have been implemented in this respect, assessing either the data aspect, the ontology schema or both. Unlike in Linked Data Quality [27] and Data Profiling [1], there is still a lack of concentrated effort to consolidate the various approaches and methods taken by different researchers to identify and obtain a subset of metrics that best represent the quality of ontologies. More effort is also needed to design tools that help ontology engineers, data producers and data publishers, not only to obtain metric measures, but also provide valuable insights into possible lack of quality in the ontologies under test. Visualisation tools have so far been mainly used to obtain a visual representation of ontologies, but not as an alternative way to visualise quality aspects.

The main objectives and contributions of this paper are the following:

Objective 1: *Identify and survey existing ontology and data quality metrics*
Contribution 1: This will be achieved through a systematic review of existing literature on quality metrics that have been used in various research fields including ontologies, database schemas, XML schemas, object-oriented designs, software engineering and hierarchical designs in general.

Objective 2: *Investigate frameworks and tools that enable the quality assessment of ontologies and visualise different quality aspects*
Contribution 2: In this article we will propose a preliminary framework that merges two known Linked Data tools with regard to data quality and ontology visualisation, in order to enable the visualisation of ontology quality.

The remaining sections of this paper are organised as follows: Section 2 presents the methodology and initial results of the survey to identify important metrics. The section shows how metrics are classified according to the categories and dimensions pertaining to the ISO Standard 25012 for Data Quality. Section 3 discusses and reviews existing visualisation tools and proposes an alternative way of looking at the quality of ontologies through the use of visualisation techniques.

2 Classifying Quality Metrics for Ontologies

Various metrics have been proposed in recent years, some of which are now widely accepted and implemented in a number of frameworks and tools, such as those in OQuaRE [7], OntoQualitas [23] and OntoQA [25]. Yang, Z. et al. [26] describe how the quality of an ontology should be managed and evaluated in terms of its engineering and visualisation. The authors describe how quality metrics help engineers in their ontology design, thus:

(1) expected to lessen the need for maintenance and,
(2) provide means to find the most fit-for-use ontologies.

2.1 ISO/IEC 25012 Data Quality Standard

The ISO/IEC 25012 [12] is an approved standard, forming part of a series of International Standards for Software Product Quality Requirements and Evaluation (SQuaRE). The model has been adopted in various areas such as software engineering [9], ontologies [6] and to data on the World Wide Web and applications [22], to define quality measures and perform quality evaluations. It categorises fifteen quality dimensions into three main categories. We aim to classify the metrics using this standard as in ontologies we are interested in both the inherent category (such as detecting inconsistencies), as well as the system category (such as detecting dereferenceability).

2.2 Survey Methodology

In order to ensure that research is thorough and fair, a systematic review was deemed necessary. The review was carried out according to the methods mentioned in [14].

Search Strategy: Based on the objective of surveying quality metrics from different research areas, several search terms that were deemed to be more appropriate for this systematic review, were used. These included:

data quality, assessment, evaluation, linked data, ontology quality, quality metrics, software quality metrics, database quality metrics.

Repositories: The following three repositories were considered in the survey:

- ScienceDirect
- IEEE Xplore Digital Library
- ACM Digital Library

2.3 Metrics Survey

An exercise was carried out to map the metrics identified in the survey, to a category and dimension of the ISO/IEC 25012 Data Quality Standard. The standard identifies three categories, as follows:

The Inherent Category caters for metrics that measure the degree to which the model itself has quality characteristics of intrinsic nature to satisfy 'fitness for use'. This includes domain values, relationships and other metadata. In our work, we refer to the accuracy, completeness, consistency and currentness dimensions of this category. The System Category refers to quality metrics that measure the degree to which quality is maintained when the system is under specific use, and includes availability, reliability and portability. The Inherent-System Category includes dimensions that look at both Inherent and System aspects, such as compliance and understandability, to which we make reference in our work.

Table 1 to Table 7 show the metrics in their respective dimensions. Some metrics may belong to multiple dimensions or categories, however, we categorise the metrics into the most appropriate dimension.

Inherent Category Metrics Table 1 to Table 4 show the association of the metrics to the ISO 25012 Inherent Category. For example **IA** refers to the association between the Inherent Category and the Accuracy dimension.

Table 1. Accuracy Dimension

Ref.	Metric	Dimension	Reference
IA1	Incorrect Relationship	Accuracy	[20], [21]
IA2	Merging of Different Concepts in same Class	Accuracy	[21]
IA3	Hierarchy Overspecialisation	Accuracy	[21], [3]
IA4	Using a Miscellaneous Class	Accuracy	[21]
IA5	Chain of Inheritance	Accuracy	[3]
IA6	Class Precision	Accuracy	[23]
IA7	Number of Deprecated Classes and Properties	Accuracy	[11]

IA1: Incorrect Relationship: An incorrect relationship typically occurs with the vague use of 'is', instead of 'subClassOf', 'type' or 'sameAs'. As mentioned in [20], the correct use of the type of relationship is required to accurately represent the domain. As explained by [21], the relationship 'rdfs:subClassOf' is reserved for subclass relationship, 'rdf:type' for objects that belong to a particular class, and 'owl:sameAs' is used to indicate that two instances are equivalent.

IA2: Merging of Different Concepts in same Class: Every different concept should be in its own class. The anomaly occurs when two different concepts are put in the same class.

IA3: Hierarchy Overspecialisation: Overspecialisation occurs when a leaf class of an ontology (a class that is not a superclass of some other classes) does not have any instances associated with it.

IA4: Using a Miscellaneous Class: A class within the hierarchy of the ontology which is simply used to represent instances that do not belong to any of its siblings. For instance, having the class 'Fruit' with subclasses 'Orange', 'Apple', 'Pear' and 'Miscellaneous'. The 'Miscellaneous' class might simply be capturing the rest of the fruits, without any distinction between them, thereby lacking accuracy.

IA5: Chain of Inheritance: An undesirable inheritance chain may occur when a large part of an ontology exists where each class in the chain has only one subclass (for example a section of the ontology with a chain of six classes, each of which has only one subclass and has no siblings). This might mean that some aggregation of the concepts defined in that section might be required.

IA6: Class Precision: This metric is calculated over a given frame of reference (existing resources or sources of data with which the ontology may be evaluated) and tests precision of the ontology. It is defined as the cardinality of the intersection between classes in the ontology and classes in the frame, divided by the total number of classes in the ontology. Effectively this is a percentage of the number of classes common between the ontology and the test data source, with respect to the total number of classes in the ontology. For example, assuming an ontology of fifty classes, of which, forty are present in the test data source, the ontology precision would be 80%. There is 20% of the ontology which is not relevant to the test data source.

IA7: Number of Deprecated Classes and Properties: This metric addresses parts of an ontology which are marked as deprecated, identified by 'owl:DeprecatedClass' or 'owl:DeprecatedProperty'. Deprecated sections are normally not updated anymore and might be superseded by newer classes or properties. This problem could either be within the ontology itself, or pointing to external references that have since been deprecated. It must be noted here that, having an ontology with a deprecated class or property is not necessarily a quality problem. In fact, in certain situations it might be desirable to leave the classes and properties within the ontology and mark them as deprecated (rather than deleting them), as there might be other ontologies that are currently referencing the deprecated elements. Deleting those elements might make the other ontologies unusuable. What we mean here is that, new ontologies developed after an element or property has been deprecated, should not ideally make use of those elements (but rather use the new elements).

Table 2. Completeness Dimension

Ref.	Metric	Dimension	Reference
IC1	Number of Isolated Elements	Completeness	[21]
IC2	Missing Domain or Range in Properties	Completeness	[21]
IC3	Class Coverage	Completeness	[23]
IC4	Relation Coverage	Completeness	[23]

IC1: Number of Isolated Elements: Elements, including classes, properties and datatypes are considered isolated if they do not have any relation to the rest of the ontology (declared but not used).

IC2: Missing Domain or Range in Properties: Properties should be accompanied by their domain and range. Missing information about the properties may cause lack of completeness and may result in less accuracy and more inconsistencies. This does not always and necessarily indicate a quality problem. There might be cases, for instance in Linked Data, where it is desirable for a property to be open (not being bound to a particular domain or specific range).

IC3: Class Coverage: This metric is calculated over a given frame of reference and determines the amount of coverage of a given ontology. It is defined as the cardinality of the intersection between classes in the ontology and classes in the frame, divided by the total number of classes in frame. Effectively this is a percentage of the number of classes common between the ontology and the test data source, with respect to the total number of classes in the test data source.

For example, assuming a test data source of sixty classes, of which, forty are present in the ontology, the ontology coverage would be 67%. There is 33% of the test data source which is not covered by the ontology.

IC4: Relation Coverage: This is similar to class coverage, but is defined as the cardinality of the intersection between relations in the ontology and relations in the frame, divided by the total number of relations in frame.

Table 3. Consistency Dimension

Ref.	Metric	Dimension	Reference
IO1	Number of Polysemous Elements	Consistency	[21]
IO2	Including Cycles in a Class Hierarchy	Consistency	[20],[10],[21]
IO3	Missing Disjointness	Consistency	[20],[10],[21]
IO4	Defining Multiple Domains/Ranges	Consistency	[21]
IO5	Creating a Property Chain with One Property	Consistency	[21]
IO6	Lonely Disjoints	Consistency	[3]
IO7	Tangledness (two methods)	Consistency	[7]
IO8	Semantically Identical Classes	Consistency	[23]

IO1: Number of Polysemous Elements: Number of properties, objects or datatypes that are referred by the same identifier. A quality issue arises if, in a given ontology, there are multiple classes and/or properties which are conceptually different but have the same identifier. For example, 'man' might refer to different but related concepts, such as referring to 'the human species' or a 'male person'.

IO2: Including Cycles in a Class Hierarchy: Identified by [10] as circulatory errors, this condition typically occurs, for example, when a class C_1 is defined as a superclass of class C_2, and C_2 is defined as a superclass of C_1 at the same time. C_1 and C_2 may not necessarily be directly linked, thus cycles may form at different depths, d.

IO3: Missing Disjointness: Gomez-Perez et al. in [10] qualifies that subclasses of a class which are disjoint from each other (a subclass can only be of one type), should specify this disjointness in the ontology.

IO4: Defining Multiple Domains/Ranges: Multiple domains and ranges are allowed, however, these should not be in conflict with each other (that is, no two domains or ranges should contradict each other). A quality issue arises when multiple definitions are inconsistent.

IO5: Creating a Property Chain with One Property: This metric refers to the use of the OWL construct 'owl:propertyChainAxiom' to set a property as being composed of several other properties. The anomaly occurs when a property chain includes only one property in the compositional part. For example, declaring the property 'grandparent' as a property chain, but including only one property 'parent' within it (instead of the required two 'parent' properties).

IO6: Lonely Disjoints: As mentioned in [3], a class C is referred to as a lonely disjoint when the ontology specifies that this class is disjoint with some other classes C_A and C_B, but C is not a sibling of C_A and C_B.

IO7: Tangledness: This is defined as the mean number of classes with more than one direct ancestor. Another measure of tangledness is defined as the mean number of direct ancestor of classes with more than one direct ancestor.

IO8: Semantically Identical Classes: This anomaly occurs when an ontology includes multiple classes with the same semantics (referring to the same concept).

Table 4. Currentness Dimension

Ref.	Metric	Dimension	Reference
IU1	Freshness	Currentness	[18]

IU1: Freshness: This is defined by [18] as a measure indicating how updated a given piece of information is. The authors define a similar metric, 'newness' as a measure to indicate how data was created in a timely manner.

Inherent-System Category Metrics Table 5 and Table 6 show the association of metrics to the ISO 25012 Inherent-System Category (IS).

Table 5. Compliance Dimension

Ref.	Metric	Dimension	Reference
ISM1	No OWL Ontology Declaration	Compliance	[21]
ISM2	Ambiguous Namespace	Compliance	[21]
ISM3	Namespace Hijacking	Compliance	[21]
ISM4	Number of Syntax Errors	Compliance	[11]

ISM1: No OWL Ontology Declaration: Ontologies must ensure that the 'owl:Ontology' tag is provided, which includes meta-data specific to the ontology such as version, license and dates, and to make reference to other ontologies.

ISM2: Ambiguous Namespace: The absence of the ontology URI and the namespace 'xml:base' will cause the ontology namespace to be matched to its location. This may result in an unstable ontology which causes its namespace to change depending on its location.

ISM3: Namespace Hijacking: Hijacking occurs when an ontology makes reference to terms T, properties P or objects O from another namespace K, where that namespace K does not really have any definitions for T, P and O.

ISM4: Number of Syntax Errors: This is a running total of the number of syntax errors found in a given ontology.

Table 6. Understandability Dimension

Ref.	Metric	Dimension	Reference
ISU1	Missing Annotations	Understandability	[21]
ISU2	Property Clumps	Understandability	[3]
ISU3	Using Different Naming Conventions	Consistency	[21]

ISU1: Missing Annotations: Elements of an ontology should have human readable annotations that label them, such as the use of 'rdfs:label' or the label 'skos:prefLabel'.

ISU2: Property Clumps: Clumps occur when a collection of elements (properties, objects) are included as a group in a number of class definitions. In such cases, [3] argue that the ontology may be improved by defining an abstract concept as an aggregation of the clump. A trivial example would be the common use of properties 'house', 'street', 'town' and 'country', together in different places within an ontology. An abstract single concept 'address' may be defined to include such properties.

ISU3: Using Different Naming Conventions: This is an inconsistency in the way concepts, classes, properties and datatypes are written.

System Category Metrics Table 7 shows the association of metrics to the ISO 25012 System Category (S).

Table 7. Availability Dimension

Ref.	Metric	Dimension	Reference
SA1	Dereferenceability	Availability	[21]

SA1: Dereferenceability: This indicates whether a given ontology is readily available online.

3 Visualisation

3.1 Visualising Ontologies

Various attempts have been made at visualising ontologies, mostly representing them as graphs which depict the way concepts are connected together. Typically, these attempts render force-directed hierarchical structures that present a nice, intuitive and useful way of displaying ontologies. Lohmann, S. et al. [16] argue that most visualisations lack in some respect. Some implementations such as OWLViz[3] and OntoTrack [15] just present the user with the hierarchy of concepts. Other systems provide more detail but lack in aspects such as datatypes and characteristics that are necessary to better understand what ontologies are really representing. These include systems such as OntoGraf[4] and FlexViz [8]. The authors further argue that VOWL is built with a comprehensive language for representation and visualisation of ontologies which can be understood by both engineers with expertise in ontologies and design, as well as by others who may be less-knowledgeable in the area. Their implementation is designed for the Web Ontology Language, OWL. This, along with the fact that VOWL is released under the MIT license and is fully available and extensible enough, is main reason why it is being used in this work to study how visualisation techniques may help ontology engineers and users to assess quality.

3.2 Visualising Ontology Quality - A Preliminary Investigation in Building a Pipeline between Luzzu and VOWL

In order to tackle Objective 2, we try to merge efforts done in Linked Data quality assessment frameworks and ontology visualisation tools. In order to achieve this, we plan to investigate the outcomes of Luzzu [5], and re-use its interoperable quality results and problem reports within VOWL [16], in a proposed system (work in progress) as shown in Figure 1.

Luzzu was selected since it is a generic assessment framework, allowing for the custom definition of quality metrics. Furthermore, the output generated by Luzzu following the quality assessment, is interoperable - in the sense that we can use the same schemas Luzzu uses to output the problem report and quality metadata, in order to visualise ontology quality in VOWL. Our aim is to create an additional layer on top of VOWL to visualise ontology quality and identify quality weaknesses, as shown in Figure 2.

Areas of interest among concepts and properties are calculated according to the number of different metrics, the different groups and the nature of the metrics that fail. Different methods and visualisation techniques will be studied to determine how these can help ontology engineers and users to visualise quality

[3] http://protegewiki.stanford.edu/wiki/OWLViz
[4] http://protegewiki.stanford.edu/wiki/OntoGraf

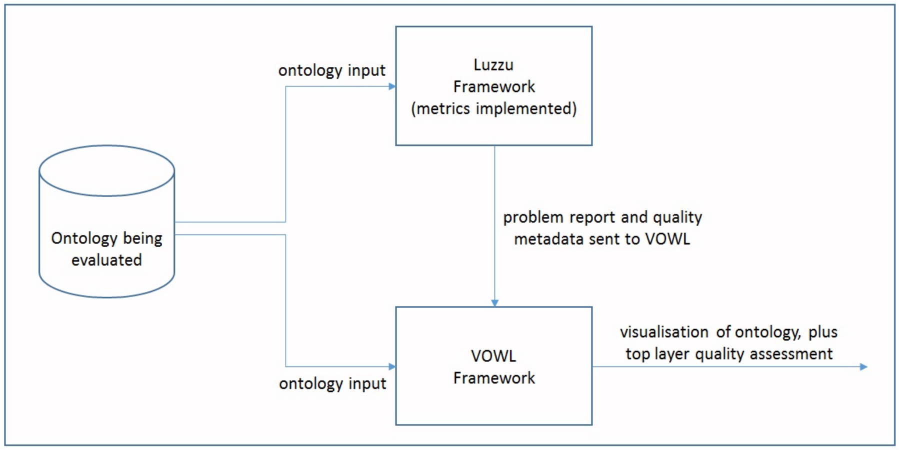

Fig. 1. Proposed System

problems as clearly as possible in such a way that they could be easily understood and interpreted correctly. The system would provide information about which metrics were used in the assessment, in such a way that it would be possible to compare two visualised quality assessments with different metrics and evaluate the effect on the given ontology.

Figure 2 shows an ontology which has been subjected to analysis. The three areas identified (highlighted) represent locations of the ontology which failed one or more tests. In this particular example, concept C_5 failed a number of tests represented here by the overlap of the three highlighted groups. An interpretation of this could be that concept C_5 might require immediate attention since it has a higher degree of weakness.

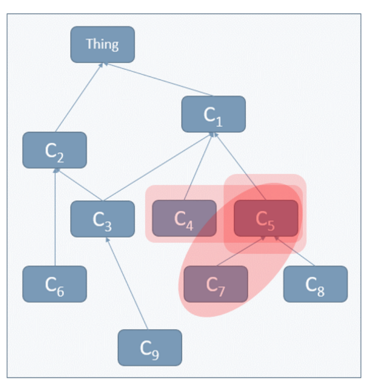

Fig. 2. Projecting Metric Information onto the Visualised Ontology

4 Final Remarks and Future Work

Ontological quality is desirable given the popularity and the important role of ontologies in communication and sharing of information across systems. This work aims at providing a comprehensive view of quality metrics for ontologies. It also looks at how visualisations can help in this process. An attempt to answer these questions is made through a survey of existing metrics from literature, obtained from different areas of computing. Correlation tests will be performed to determine sets of metrics that address the same aspects of quality. The results of the survey and correlation tests will help in identifying metrics that will then be implemented in the Luzzu framework. Ontologies are assessed using this framework, and its quality metadata and problem reports are fed into the VOWL framework, whereby an additional layer will be implemented to provide a visualisation of the quality assessment for the given ontology. As a result, we aim to provide an alternative and more intuitive way of looking at the level of quality in an ontology, achieved through visualisation techniques.

References

1. Abedjan, Z., Golab, L., Naumann, F.: Profiling relational data: a survey. The VLDB Journal. 24, 557-581 (2015).
2. Ashburner, M., Ball, C., Blake, J., Botstein, D., Butler, H., Cherry, J., Davis, A., Dolinski, K., Dwight, S., Eppig, J., Harris, M., Hill, D., Issel-Tarver, L., Kasarskis, A., Lewis, S., Matese, J., Richardson, J., Ringwald, M., Rubin, G., Sherlock, G.: Gene Ontology: tool for the unification of biology. Nature Genetics. 25, 25-29 (2000).
3. Baumeister, J., Seipel, D.: Smelly owls - Design anomalies in ontologies. Proceedings of the Eighteenth International Florida Artificial Intelligence Research Society Conference, FLAIRS 2005 - Recent Advances in Artifical Intelligence. 215-220 (2005).
4. Besbes, G., Baazaoui-Zghal, H.: Modular ontologies and CBR-based hybrid system for web information retrieval. Multimedia Tools and Applications. 74, 8053-8077 (2014).
5. Debattista, J., Auer, S., Lange, C.: Luzzu - A Methodology and Framework for Linked Data Quality Assessment. Journal of Data and Information Quality. 8, 1-32 (2016).
6. Duque-Ramos, A., Boeker, M., Jansen, L., Schulz, S., Iniesta, M., Fernndez-Breis, J.: Evaluating the Good Ontology Design Guideline (GoodOD) with the Ontology Quality Requirements and Evaluation Method and Metrics (OQuaRE). PLoS ONE. 9, e104463 (2014).
7. Duque-Ramos, A., Fernndez-Breis, J., Iniesta, M., Dumontier, M., Egaa Aranguren, M., Schulz, S., Aussenac-Gilles, N., Stevens, R.: Evaluation of the OQuaRE framework for ontology quality. Expert Systems with Applications. 40, 2696-2703 (2013).
8. Falconer, S.M., Callendar, C., Storey, M.: FLEXVIZ: Visualizing Biomedical Ontologies on the Web. International Conference on Biomedical Ontology, Software Demonstration, Buffalo, NY. 0-1, (2009).

9. Febrero, F., Calero, C., Angeles Moraga, M.: Software reliability modeling based on ISO/IEC SQuaRE. Information and Software Technology. 70, 18-29 (2016).
10. Gomez-Perez, A., Fernandez-Lopez, M., Corcho, O.: Ontological engineering. Springer, London. (2010).
11. Hogan, A., Harth, A., Passant, A., Decker, S., Polleres, A.: Weaving the pedantic Web. CEUR Workshop Proceedings. 628 (2010).
12. ISO, ISO/IEC 25012:2008Software engineering. Software product quality requirements and evaluation (SQuaRE). Data quality model, Report, International Organization for Standarization. (2009).
13. Juran, J., Godfrey, A.: Juran's Quality Handbook (5th Edition). McGraw-Hill Professional Publishing, New York, USA. (1998).
14. Kitchenham, B.: Procedures for performing systematic reviews. Technical report, Joint Technical Report Keele University Technical Report TR/SE-0401 and NICTA Technical Report 0400011T.1 (2004).
15. Liebig, T., Noppens, O.: OntoTrack - A New Ontology Authoring Approach. 4 (2004).
16. Lohmann, S., Negru, S., Haag, F., Ertl, T.: Visualizing ontologies with VOWL. Semantic Web. 7, 399-419 (2016).
17. McCray, A.: An Upper-Level Ontology for the Biomedical Domain. Comparative and Functional Genomics. 4, 80-84 (2003).
18. Mendes, C.P.N., Bizer, C., Miklos, Z., Calbimonte, J., Moraru, A., Flouris, G.: PlanetData D2.1 Conceptual model and best practices for high-quality metadata publishing.
19. Miranda, S., Orciuoli, F., Sampson, D.: A SKOS-based framework for Subject Ontologies to improve learning experiences. Computers in Human Behavior. 61, 609-621 (2016).
20. Noy, N.F., McGuinness, D.L.: Ontology Development 101: A Guide to Creating Your First Ontology. Stanford Knowledge Systems Laboratory. 25 (2001).
21. Poveda-Villaln, M., Gomez-Perez, A., Suarez-Figueroa, M.: OOPS! (OntOlogy Pitfall Scanner!):. International Journal on Semantic Web and Information Systems. 10, 7-34 (2014).
22. Rafique I., Lew P., Qanber Abbasi M., Li, Z.: Information Quality Evaluation Framework: Extending ISO 25012 Data Quality Model, World Academy of Science, Engineering and Technology - International Journal of Computer, Electrical, Automation, Control and Information Engineering. 6, 568-573 (2012).
23. Rico, M., Caliusco, M., Chiotti, O., Galli, M.: OntoQualitas: A framework for ontology quality assessment in information interchanges between heterogeneous systems. Computers in Industry. 65, 1291-1300 (2014).
24. Srinivasan, K., Devi, T.: A Comprehensive Review and Analysis on Object-Oriented Software Metrics in Software Measurement. International Journal on Computer Science and Engineering. 6, 7, 247-261 (2014).
25. Tartir, S., Arpinar, I.B., Moore, M., Sheth, A.P., Aleman-Meza, B.: OntoQA: Metric-Based Ontology Quality Analysis. IEEE Workshop on Knowledge Acquisition from Distributed, Autonomous, Semantically Heterogeneous Data and Knowledge Sources. 45-53 (2005).
26. Yang, Z., Zhang, D., Ye, C.: Evaluation metrics for ontology complexity and evolution analysis. Proceedings - IEEE International Conference on e-Business Engineering, ICEBE 2006. 162-169 (2006).
27. Zaveri, A., Rula, A., Maurino, A., Pietrobon, R., Lehmann, J., Auer, S.: Quality assessment for Linked Data: A Survey. Semantic Web. 7, 63-93 (2015).