

Towards Vagueness-Oriented Quality Assessment of Ontologies

Panos Alexopoulos¹ and Phivos Mylonas²

¹ iSOCO, Avda del Partenon 16-18, 28042, Madrid, Spain
palexopoulos@isoco.com

² Ionian University, Department of Informatics, 7 Tsirigoti Square, 49100, Corfu, Greece
fmylonas@ionio.gr

Abstract. Ontology evaluation has been recognized for a long time now as an important part of the ontology development lifecycle, and several methods, processes and metrics have been developed for that purpose. Nevertheless, vagueness is a quality dimension that has been neglected from most current approaches. Vagueness is a common human knowledge and linguistic phenomenon, typically manifested by terms and concepts that lack clear applicability conditions and boundaries such as *high*, *expert*, *bad*, *near* etc. As such, the existence of vague terminology in an ontology may hamper the latter’s quality, primarily in terms of shareability and meaning explicitness. With that in mind, in this short paper we argue for the need of including vagueness in the ontology evaluation activity and propose a set of metrics to be used towards that goal.

1 Introduction

Ontologies are formal conceptualizations of domains, describing the meaning of domain-specific data in a common, machine-processable form by means of concepts and their interrelations [6]. Their primary function is to ensure that the meaning of data exchanged between and within systems is consistent and shared, both by humans and machines. Nevertheless, as a given ontology might not always be able to perform this function in a satisfactory way for the domain and/or application scenario at hand, a thorough evaluation of it should always precede its release and usage.

Several ontology evaluation approaches have been proposed during the past years. A typical categorization scheme for these approaches regards the process by which evaluation is performed. Thus, one may find approaches that compare the ontology to some “Gold Standard” ontology and measure the former’s fitting degree to the latter [4], or approaches that do the same with data sources that are relevant to the domain an ontology is supposed to cover [5]. There are also evaluation approaches that use an ontology in some concrete application context and evaluate it along with the application [9] and, finally, approaches that have human users judging how well an ontology satisfies a set of predefined criteria and requirements [11]. At the same time, different evaluation approaches define

target different ontology dimensions. Thus, there are criteria and metrics related to the *content* of the ontology (e.g. what entities does it contain), its *structure* (e.g. how are entities related), its *usability* (e.g. how well documented it is) and its *functionality* (e.g. how well can it answer particular queries).

In this paper, we are interested in an ontology evaluation dimension, which, to the best of our knowledge, has not been considered in any existing evaluation framework, namely *vagueness*. Vagueness is a common human knowledge and language phenomenon, typically manifested by terms and concepts like *high*, *expert*, *bad*, *near* etc., and related to our inability to precisely determine the extensions of such concepts in certain domains and contexts. That is because vague concepts have typically blurred boundaries which do not allow for a sharp distinction between the entities that fall within their extension and those that do not [8]. For example, some people are borderline tall: not clearly “*tall*” and not clearly “*not tall*”.

In a previous paper, we have argued that the existence of vague terminology in ontologies and semantic data can affect in a negative way ontology comprehensibility and shareability and limit their value as a reusable source of knowledge [1]. The reason is the subjective interpretation of vague definitions that can cause **disagreements** among the people who develop, maintain or, most importantly, use a vague ontology. With that in mind, in this paper we focus on the problem of assessing if and to what extent vagueness is present in an ontology, as well as to what extent it hampers its quality, mainly in terms of meaning explicitness and shareability. To facilitate this assessment we define a set of vagueness-oriented ontology evaluation metrics and provide some guidelines for measuring them.

The structure of the rest of the paper is as follows. In the next section we explain in more detail the motivation behind our work while in section 3 we describe a set of vagueness-specific ontology evaluation criteria and metrics as well as guidelines for their assessment/measurement. Section 4 describes the process and results of applying the framework for the evaluation of actual ontologies and, finally, section 5 summarizes our work and outlines its future directions.

2 Motivation

The existence of vague terminology in ontologies and semantic data has already been identified by the community [3] [1], as well as the potential negative effects it may have on ontology comprehensibility and shareability [1]. The reason is the subjective interpretation of vague definitions that can cause **disagreements** among the people who develop, maintain or, most importantly, use a vague ontology.

Such a situation, for example, arose in a real life application scenario where, while trying to develop an electricity market ontology for an energy-related electronic marketplace, we faced significant difficulties in defining concepts like “*Critical System Process*”, “*Strategic Market Participant*” or “*High Profit Margin*”. The reason was that when, for example, we asked our domain experts to provide exemplar instances of critical processes, there were certain processes for

which there was a dispute among the experts about whether they should be regarded as critical or not. As it turned out, the source of the problem was not only that different domain experts had different criteria of process criticality, but also that no one could really decide which of those criteria were sufficient to classify a process as critical. In other words, the problem was the vagueness of the predicate “*critical*”, manifested by the existence of borderline cases regarding its applicability.

More generally, as we have shown in [1], vagueness in ontologies can cause problems in scenarios involving i) the structuring of data with a vague ontology (where disagreements among experts on the validity of vague statements may occur), ii) the utilization of vague facts in ontology-based systems (where reasoning results might not meet users’ expectations) and iii) the integration of vague semantic information (where the merging of particular vague elements can lead to semantic data that will not be valid for all its users).

3 Vagueness Metrics for Ontologies

3.1 Vagueness Spread

In an ontology the elements that may be vague are typically concepts, relations, attributes and datatypes. A concept is vague if, in the given domain, context or application scenario, it admits borderline cases, namely if there are (or could be) individuals for which it is indeterminate whether they instantiate the concept. Primary candidates for being vague are concepts that denote some phase or state (e.g., *adult*, *child*) as well as attributions, namely concepts that reflect qualitative states of entities (e.g., *red*, *big*, *broken* etc.).

Similarly, a relation is vague if there are (or could be) pairs of individuals for which it is indeterminate whether they stand in the relation. The same applies for attributes and pairs of individuals and literal values. Finally, a vague datatype consists of a set of vague terms which may be used within the ontology as attribute values. For example, the attribute *performance*, which normally takes as values integer numbers, may also take as values terms like *very poor*, *poor*, *mediocre*, *good* and *excellent*. Primary candidates for generating such terms are gradable attributes such as size or height which give rise to terms such as *large*, *tall*, *short* etc.

Given the above, the first vagueness-related metric we propose for an ontology is **Vagueness Spread (VS)**, namely the ratio of the number of ontological elements (classes, relations and datatypes) that are vague (VOE), divided by the total number of elements (OE):

$$VS = \frac{|VOE|}{|OE|} \quad (1)$$

This metric practically reflects the extent to which vagueness is present in the ontology and it provides an indication of the ontology’s potential comprehensibility and shareability; an ontology with a high value of vagueness spread is less explicit and shareable than an ontology with a low value.

Calculation of this metric is typically performed in a manual fashion by domain experts who need to identify the elements of the ontology that are vague. This can be done by analyzing the elements’ definitions and deciding whether they admit borderline cases. Thus, for example, the definition of the ontology class “*StrategicClient*” as “*A client that has a high value for the company*” is vague while the definition of “*AmericanCompany*” as “*A company that has legal status in the Unites States*” is not. As part of our ongoing work, we are in the process of developing a vague sense classifier that may be help domain experts detect vague ontological definitions in a semi-automatic fashion; nevertheless this is to be reported in a future work.

3.2 Vagueness Explicitness

A vague ontology element that is explicitly identified and documented as such, is more usable than one that’s not. The reason is that without such an explicit characterization, ontology users may not realize the fact that the element is vague and use it, thus risking the potential consequences mentioned in the introduction.

Therefore, a second metric we propose is **Vagueness Explicitness (VE)**, namely the ratio of the number of vague ontological elements that are explicitly identified as such (EVOE), divided by the total number of vague elements (VOE):

$$VE = \frac{|EVOE|}{|VOE|} \quad (2)$$

Obviously, the higher the value of this metric is, the better is the ontology. As far as its calculation is concerned, this is done also in a manual fashion by checking the elements’ definitions.

3.3 Vagueness Intensity

As suggested in the introduction, vague ontology elements can be problematic because of the disagreement they may cause among the ontology’s users. The higher this disagreement is, the more problems the element is likely to cause. Thus, another vagueness-related metric we define is **Vagueness Intensity (VI)**, namely the degree to which the ontology’s users disagree on the validity of the (potential) instances of the ontology elements. The exact formula for this metric depends on the way one decides to measure user agreement. Our proposed approach is to i) consider a sample set of vague element instances, ii) ask from a group of potential ontology users to denote whether and to what extent they believe these instances are valid and iii) Measure the inter-agreement between users using Cohen’s Kappa (if users simply say “Agree” or “Disagree”) or its weighted version (if users rate their agreement in some scale) [10]. For example, the value of kappa for judges who either agree or disagree to a vague statement’s validity is given as follows:

$$\kappa = \frac{Pr(\alpha) - Pr(e)}{1 - Pr(e)} \quad (3)$$

where $Pr(\alpha)$ is the proportion of times the judges agree and $Pr(e)$ is the proportion of times they are expected to agree by chance alone. Complete agreement corresponds to $\kappa = 1$, and lack of agreement (i.e. purely random coincidences of rates) to $\kappa = 0$. Then, the value of VI for a set of S vague statements is the average of the disagreements for each statement, namely:

$$VI = \frac{1}{|S|} \sum_{s \in S} (1 - \kappa_s) \quad (4)$$

4 Vagueness Evaluation Examples

To illustrate how our proposed metrics can be applied in practice, we applied them to two concrete ontologies. The first was the Citation Typing Ontology (or CiTO³), a publicly available ontology that enables characterization of the nature or type of citations, both factually and rhetorically, permitting these descriptions to be published on the Web. CiTO consists primarily of relations, many of which are vague (e.g., relations *plagiarizes* and *citesAsAuthority*). In order to measure the vagueness spread of the particular ontology we had two domain experts manually identify the elements that were vague. In the end, we got 27 vague elements and 17 non-vague, resulting in a VS of 0.61. Vagueness explicitness, in turn, was 0, practically because none of the vague elements were identified as such. Finally, vagueness intensity is derived from a CiTO usage experiment [7] where the authors measured the kappa value for 5 raters over 104 relation instances and obtained a kappa value of 0.16.

The second ontology was an enterprise ontology, developed in the context of a decision support system for tender call evaluation [2]. This ontology defined basic concepts and relations regarding an enterprise’s internal and external environment (e.g. *Business Function*, *Employee*, *Client*, etc.), several of which were vague (e.g. *Competitor*, *High Potential Employee*, etc.). Vagueness spread in this case was lower than CiTO, (0.25) but its intensity (which we measured by having members of the enterprise assess the validity of sample instances of the vague elements) was significant (0.75). Finally, vagueness explicitness was 1.0, practically because we had developed the ontology ourselves, making sure that vagueness was properly documented.

Table 1. Correspondence of Metamodel Elements to Required Vagueness Aspects

Ontology	VS	VE	VI
CiTO	0.61	0.0	0.84
Enterprise Ontology	0.25	1.0	0.75

³ <http://purl.org/spar/cito/>

5 Conclusions and Future Work

In this work we considered the phenomenon of vagueness and discussed its role in the ontology evaluation process. We introduced three novel metrics, i.e., Vagueness Spread (VS), Vagueness Explicitness (VE) and Vagueness Intensity (VI) that may be used in recognizing and calculating vagueness early enough within an ontology structure, thus aiding towards its quality optimization. Among our future work is to automate the vague element identification process as well as to apply the metrics to a larger number of publicly available ontologies and semantic datasets.

References

1. P. Alexopoulos, B. Villazon-Terrazas, and J. Z. Pan. Towards vagueness-aware semantic data. In *URSW*, volume 1073 of *CEUR Workshop Proceedings*, pages 40–45. CEUR-WS.org, 2013.
2. P. Alexopoulos, M. Wallace, K. Kafentzis, and A. Thomopoulos. A fuzzy knowledge-based decision support system for tender call evaluation. In *AIAI*, volume 296 of *IFIP Advances in Information and Communication Technology*, pages 51–59. Springer, 2009.
3. F. Bobillo and U. Straccia. Fuzzy ontology representation using owl 2. *International Journal of Approximate Reasoning*, 52(7):1073–1094, Oct. 2011.
4. J. Brank, D. Madenic, and M. Groblenik. Gold standard based ontology evaluation using instance assignment. In *Proceedings of the 4th Workshop on Evaluating Ontologies for the Web (EON2006)*, Edinburgh, Scotland, May 2006.
5. C. Brewster, H. Alani, S. Dasmahapatra, and Y. Wilks. Data-driven ontology evaluation. In *Proceedings of the Language Resources and Evaluation Conference (LREC 2004)*, pages 164–168, Lisbon, Portugal, 2004. European Language Resources Association.
6. B. Chandrasekaran, J. Josephson, and R. Benjamins. What are ontologies and why do we need them? *IEEE Intelligent Systems*, 14(1):Page 20–26, January - February 1999.
7. P. Ciancarini, A. D. Iorio, A. G. Nuzzolese, S. Peroni, and F. Vitali. Characterising citations in scholarly articles: An experiment. In *AIC@AI*IA*, volume 1100 of *CEUR Workshop Proceedings*, pages 124–129. CEUR-WS.org, 2013.
8. D. Hyde. *Vagueness, Logic and Ontology*. Ashgate New Critical Thinking in Philosophy, 2008.
9. R. Porzel and R. Malaka. A task-based approach for ontology evaluation. In *Proceedings of ECAI 2004 Workshop on Ontology Learning and Population*, Valencia, Spain, August 2004.
10. J. Sim and C. C. Wright. The kappa statistic in reliability studies: Use, interpretation, and sample size requirements. *Physical Therapy*, March 2005.
11. S. Tartir, I. B. Arpinar, M. Moore, A. P. Sheth, and B. Aleman-Meza. OntoQA: Metric-based ontology quality analysis. In *Proceedings of IEEE Workshop on Knowledge Acquisition from Distributed, Autonomous, Semantically Heterogeneous Data and Knowledge Sources*, 2005.