

The Current Landscape of Pitfalls in Ontologies

C. Maria Keet¹, Mari Carmen Suárez-Figueroa² and María Poveda-Villalón²

¹*School of Mathematics, Statistics, and Computer Science, University of KwaZulu-Natal, and UKZN/CSIR-Meraka Centre for Artificial Intelligence Research, Durban, South Africa*

²*Ontology Engineering Group, Departamento de Inteligencia Artificial, Universidad Politécnica de Madrid, Madrid, Spain*
keet@ukzn.ac.za, {mcsuarez,mpoveda}@fi.upm.es

Keywords: Ontology Development : Ontology Quality : Pitfall

Abstract: A growing number of ontologies are already available thanks to development initiatives in many different fields. In such ontology developments, developers must tackle a wide range of difficulties and handicaps, which can result in the appearance of anomalies in the resulting ontologies. Therefore, ontology evaluation plays a key role in ontology development projects. OOPS! is an on-line tool that automatically detects pitfalls, considered as potential errors or problems, and thus may help ontology developers to improve their ontologies. To gain insight in the existence of pitfalls and to assess whether there are differences among ontologies developed by novices, a random set of already scanned ontologies, and existing well-known ones, data of 406 OWL ontologies were analysed on OOPS!'s 21 pitfalls, of which 24 ontologies were also examined manually on the detected pitfalls. The various analyses performed show only minor differences between the three sets of ontologies, therewith providing a general landscape of pitfalls in ontologies.

1 INTRODUCTION

A growing number of ontologies are already available in different domains thanks to ontology development initiatives and projects. However, the development of ontologies is not trivial. Early ontology authoring suggestions were made by (Noy and McGuinness, 2001), and (Rector et al., 2004) present the most common problems, errors, and misconceptions of understanding OWL DL based on their experiences teaching OWL. OWL 2 DL contains more features and there is a much wider uptake of ontology development by a more diverse group of modellers since. This situation increases the need for training, for converting past mistakes into useful knowledge for ontology authoring, and it requires a clear notion of ontology quality both in the negative sense (what are the mistakes?) and in the positive (when is some representation good?). Several steps have been taken with respect to quality in the negative sense, such as to identify antipatterns (Roussey et al., 2009) and to create a catalogue of common pitfalls—understood as potential errors, modelling flaws, and missing good-practices in ontology development—in OWL ontologies (Poveda et al., 2010; Poveda-Villalón et al., 2012), and in the positive sense by defining good and ‘safe’ object property expressions (Keet, 2012) and taxonomies (Guarino and Welty, 2009). The cata-

logue of common pitfalls included 29 types of pitfalls at the time of evaluation and 21 of them are automatically detected by the online Ontology Pitfall Scanner! (OOPS! <http://www.oeg-upm.net/oops>). With the automation of scanning pitfalls as well as advances in ontology metrics, this now provides the opportunity to obtain quantitative results, which has been identified as a gap in the understanding of ontology quality before (Vrandečić, 2009). Here, we are interested in answering two general questions, being:

- A. What is the prevalence of each of those pitfalls in existing ontologies?
- B. To what extent do the pitfalls say something about quality of an ontology?

The second question can be broken down into several more detailed questions and hypotheses, which one will be able to answer and validate or falsify through a predominantly quantitative analysis of the ontologies:

1. Which anomalies that appear in OWL ontologies are the most common?
2. Are the ontologies developed by experienced developers and/or well-known or mature ontologies ‘better’ in some modelling quality sense than the ontologies developed by novices? This is refined into the following hypotheses:
 - (i) The prevalence and average of pitfalls is significantly higher in ontologies developed by novices compared to ontologies deemed estab-

lished/mature.

- (ii) The kind of pitfalls observed in novices' ontologies differs significantly from those in well-known or mature ontologies.
- (iii) The statistics on observed pitfalls of a random set of ontologies is closer to those of novices' ontologies than the well-known or mature ones.
- (iv) There exists a positive correlation between the detected pitfalls and the size or number of particular elements of the ontology.
- (v) There exists a positive correlation between the detected pitfalls and the DL fragment of the OWL ontology.

To answer these questions, we used the 362 ontologies scanned by OOPS! over the past year, 23 novices ontologies, and 21 ontologies that are generally considered to be well-known, where the latter two sets were also scanned by OOPS! and evaluated manually. Although all 21 types of pitfalls have been detected, the most common pitfalls concern lack of annotations and domain and range axioms, and issues with inverses, and to some extent creating unconnected ontology elements and using a recursive definition. The results falsify hypotheses (i), (ii), and (v), partially validate (iv)—for novices, the number of pitfalls/ontology does relate to the size and complexity of the ontology—and validate (iii); i.e., there are no striking differences between the three sets of ontologies, therewith providing a general landscape of pitfalls in ontologies.

In the remainder of this paper, we describe the state of the art in Section 2, report on the experimental evaluation in Section 3, and conclude in Section 4.

2 STATE OF THE ART

When developing ontologies, developers must tackle a wide range of difficulties, which are related to the inclusion of anomalies in the modelling. Thus, ontology evaluation, which checks the technical quality of an ontology against a frame of reference, plays a key role when developing ontologies. To help developers during the ontology modelling, early ontology authoring guidelines to avoid typical errors when modelling ontologies were provided in (Noy and McGuinness, 2001). Such guidelines help developers to prevent errors related to the definition of classes, class hierarchies, and properties during frame-based ontology developments. Rector and colleagues (Rector et al., 2004) help with the precise meaning of OWL DL and provide some guidelines on how to avoid diverse pitfalls when building OWL DL ontologies. These pitfalls were mainly related to (a) the

failure to make information explicit, (b) the mistaken use of universal and existential restrictions, (c) the open world reasoning, and (d) the effects of domain and range constraints. A classification of errors was identified during the evaluation of consistency, completeness, and conciseness of ontology taxonomies (Gómez-Pérez, 2004). First steps towards a catalogue of common pitfalls started in 2009 (Poveda et al., 2010) leading to a first stable version in (Poveda-Villalón et al., 2010). This catalogue is being maintained and is accessible on-line as part of the OOPS! portal. OOPS! (Poveda-Villalón et al., 2012) is a web-based tool for detecting potential pitfalls, currently providing mechanisms to automatically detect a subset of 21 pitfalls of those included in the catalogue and therewith helping developers during the ontology validation activity. Related to the aforementioned catalogue of pitfalls, is the identification of a set of antipatterns (Roussey et al., 2009). Theory-based methods to help developers to increase ontology quality include defining good and 'safe' object property expressions (Keet, 2012) and ontologically sound taxonomies (Guarino et al., 2009). To help developers during the ontology evaluation activity, there are different approaches: (a) comparison of the ontology to a "gold standard", (b) use of the ontology in an application and evaluation of the results, (c) comparison of the ontology with a source of data about the domain to be covered, and (d) evaluation by human experts who assess how the ontology meets the requirements (Brank et al., 2005). A summary of generic guidelines and specific techniques for ontology evaluation can be found in (Sabou and Fernandez, 2012). An ontology evaluation approach based on the three following layers is presented in (Gangemi et al., 2006): (1) O2 (a meta-ontology), (2) oQual (a pattern based on O2 for Ontology Quality), and good (for Quality-Oriented Ontology Description). This allows one to measure the quality of an ontology relative to structural, functional, and usability-related dimensions. A compendium of criteria describing good ontologies is reported in (Vrandečić, 2009) (including accuracy, adaptability, clarity, completeness, computational efficiency, conciseness, consistency/coherence and organizational fitness) and it presents a review of domain and task-independent evaluation methods related to vocabulary, syntax, structure, semantics, representation and context aspects.

To the best of our knowledge, what is missing at present in the ontology and evaluation field is a quantitative analysis of the most common pitfalls developers include in the ontologies. Based on this study, one then may create a relevant set of guidelines to help developers in the task of developing ontologies and refine ontology quality criteria.

3 EXPERIMENTAL EVALUATION OF PITFALLS IN ONTOLOGIES

3.1 Materials and Methods

3.1.1 Data collection

With the aim of identifying the most common pitfalls typically made when developing ontologies in different contexts and domains, we have collected and analyzed 44 ontologies (Set1 and Set2) and used the data stored in OOPS! for a random set (Set3):

Set1: 23 ontologies in different domains (a.o., furniture, tennis, bakery, cars, soccer, poker, birds, and plants) developed by novices. These ontologies were developed as a practical assignment by Computer Science honours (4th year) students attending the course “Ontologies & Knowledge bases (OKB718)” in 2011 and 2012 at the University of KwaZulu-Natal.

Set2: 21 existing well-known ontologies that may be deemed ‘mature’ in the sense of being a stable release, well-known, a real OWL ontology (i.e., no toy ontology nor a tutorial ontology, nor an automated thesaurus-to-OWL file), the ontology is used in multiple projects including in ontology-driven information systems, and whose developers have ample experiences in and knowledge of ontologies, and the selected ontologies are in different subject domains; a.o., DOLCE, BioTop, and GoodRelations.

Set3: 362 ontologies analyzed with OOPS! They were selected from the 614 times that ontologies were submitted between 14-11-2011 and 19-10-2012. The full set was filtered as follows: maintain those repeated ontologies for which OOPS! obtained different results in each evaluation, eliminate those repeated ontologies for which OOPS! obtained the same results in every evaluation, and eliminate those ontologies whose namespace is deferenceable but it does not refer to an ontology.

OOPS! output for the three sets, including calculations, manual analyses of OOPS! detected pitfalls for ontologies in Set1 and Set2, and the names and URIs of the ontologies of Set2 and the names of the ontologies in Set1, are available at <http://www.oeg-upm.net/oops/material/KEOD2013/pitfallsAnalysis.xlsx>.

All ontologies are evaluated by being scanned through OOPS!, which checks the ontology on most pitfalls that have been collected in the pitfall catalogue and discussed in earlier works (Poveda et al., 2010; Poveda-Villalón et al., 2012) and are taken at face value for this first quantitative evaluation:

Creating synonyms as classes (P2); Creating the relationship “is” instead of using `rdfs:subClassOf`, `rdf:type` or `owl:sameAs` (P3); Creating unconnected ontology elements (P4); Defining wrong inverse relationships (P5); Including cycles in the hierarchy (P6); Merging different concepts in the same class (P7); Missing annotations (P8); Missing disjointness (P10); Missing domain or range in properties (P11); Missing equivalent properties (P12); Missing inverse relationships (P13); Swapping intersection and union (P19); Misusing ontology annotations (P20); Using a miscellaneous class (P21); Using different naming criteria in the ontology (P22); Using recursive definition (P24); Defining a relationship inverse to itself (P25); Defining inverse relationships for a symmetric one (P26); Defining wrong equivalent relationships (P27); Defining wrong symmetric relationships (P28); and Defining wrong transitive relationships (P29). Detailed descriptions are available online from the pitfall catalogue at <http://www.oeg-upm.net/oops/catalogue.jsp>. Note that OOPS! analyses also properly imported OWL ontologies, i.e., when they are available and dereferenceable online at the URI specified in the import axiom.

In addition, we collected from the ontologies of Set1 and Set2: DL sublanguage as detected in Protégé 4.1, number of classes, object and data properties, individuals, subclass and equivalence axioms.

3.1.2 Analyses

The data was analysed by computing the following aggregates and statistics. The basic aggregates for the three sets are: (a) percentage of the incidence of a pitfall; (b) comparison of the percentages of incidence of a pitfall among the three sets; (c) average, median, and standard deviation of the pitfalls per ontology and compared among the three sets; and (d) average, median, and standard deviation of the pitfall/ontology.

For Set1 and Set2 ontologies, additional characteristics were calculated, similar to some of the ontology metrics proposed elsewhere (Vrandečić, 2009; Gangemi et al., 2006). Let $|C|$ denote the number of classes, $|OP|$ the number of object properties, $|DP|$ the number of data properties, $|I|$ the number of individuals, $|Sax|$ the number of subclass axioms, and $|Eax|$ the number of equivalences in an ontology. The number of *Ontology Elements* (OE) is computed by Eq. 1, and an approximation of the *Ontology Size* (OS) by Eq. 2.

$$OE = |C| + |OP| + |DP| + |I| \quad (1)$$

$$OS = |C| + |OP| + |DP| + |I| + |Sax| + |Eax| \quad (2)$$

We use two measures for quantifying the ‘complexity’ of the ontology. First, an *Indirect Modelling*

Complexity (IMC) is computed based on the axioms present (Eq. 3), where a lower value indicates a more complex ontology with relatively more axioms declaring properties of the classes compared to a lightweight ontology or bare taxonomy.

$$IMC = |C| : (|Sax| + |Eax|) \quad (3)$$

Second, the OWL features used are analysed twofold: (i) by calculating the overall percentage of use of \mathcal{S} , \mathcal{R} , \mathcal{O} , \mathcal{I} , \mathcal{Q} and \mathcal{D} , i.e., a rough measure of the OWL 2 DL features used; (ii) by converting the DL fragment into a numerical value, where \mathcal{AL} is given the lowest value of 0 and \mathcal{SRIOQ} the highest value of 10, to be used in correlation calculations (see below). The DL fragment and IMC are compared as well, for they need not be similar (e.g., a bare taxonomy with one object property declared reflexive already ‘merits’ detection of an \mathcal{R} , but actually is still a simple ontology with respect to the subject domain represented, and, *vv.*, an ontology can be comprehensive with respect to the subject domain, but originally developed in OWL DL but not updated since OWL 2).

Basic correlations are computed for the ontology sizes and complexities with respect to the pitfalls, and detailed correlations are computed for certain individual pitfalls: P5, P11, P13, P25, P26, P27, P28, and P29 are pitfalls specific to object properties, hence, the amount of properties in the ontologies may be correlated to the amount of pitfalls detected, and likewise for P3, P6, P7, P10, P21, and P24 for classes, and P8 for classes and ontology elements.

Finally, manual qualitative analyses with ontologies in Set1 and Set2 were conducted on possible false positives and additional pitfalls.

3.2 Results

We first present the calculations and statistics, and subsequently a representative selection of the qualitative evaluation of the ontologies in Set1 and Set2.

3.2.1 Aggregated and analysed data

The raw data of the ontologies evaluated with OOPS! are available online at <http://www.oeg-upm.net/oops/material/KEOD2013/pitfallsAnalysis.xlsx>.

The type of mistakes made by novice ontology developers are: P4, P5, P7, P8, P10, P11, P13, P19, P22, P24, P25, P26, P27, and P29. The percentages of occurrence of a pitfall over the total set of 23 ontologies in Set1 is included in Fig. 1, the average amount of pitfalls is shown in Fig. 2, and aggregate data is listed in Table 1. The analogous results for Set3 are shown in Figs. 1 and Fig. 2, and in Table 1, noting that all OOPS! pitfalls have been detected in Set3 and that

the median amount of pitfalls/ontology is similar to that of Set1. The high aggregate values are caused by a few ontologies each with around 5000 or more detected pitfalls; without P8 (annotations), there are three ontologies that have more than 1000 detected pitfalls at the time of scanning the ontology. The results obtained with the 21 well-known ontologies (Set2) can be found in the same table and figures, and include pitfalls P2, P4, P5, P7, P8, P10, P11, P12, P13, P19, P20 (0 upon manual assessment), P21, P22, P24, P25, P26, P27, and P29, noting that the percentages and averages differ little from those of the novices and random ones. The high aggregate values for Set2 is largely due to OBI with a pitfall count of 3771 for P8 (annotations) and DMOP with a pitfall count of 866 for P8; without P8, OBI, DMOP, and the Government Ontology exceeded 100 pitfalls due to P11 (missing domain and range axioms) and P13 (missing inverses—but see also below). P8 is an outlier both in prevalence and in quantity for all three sets of ontologies and only some of the ontologies have very many missing annotations, which skews the average, as can be observed from the large standard deviations.

For Set1 and Set2, we collected data about the content of the ontologies and analysed them against the pitfalls, as described in Section 3.1. The usage of the OWL 2 DL features in Set1 are: \mathcal{S} 44%, \mathcal{R} 26%, \mathcal{I} 83%, \mathcal{O} 26%, \mathcal{Q} 52%, and \mathcal{D} 17%, whereas for Set2, the percentages are 62%, 19%, 81%, 24%, 5%, and 86%, respectively; the difference is largely due to the difference in timing of the development of the ontology, with some of the well-known ontologies having been developed before the OWL2 standard, and the use of data properties was discouraged in the lectures for the ontologies in Set1. In order to include the DL fragment in the analyses, we assigned values to the fragments prior to analysis, ranging from a value of 0 for an ontology in $\mathcal{AL}(\mathcal{D})$ to 10 for an ontology in $\mathcal{SRIOQ}(\mathcal{D})$, and intermediate values for others (e.g., $\mathcal{ALCHI}(\mathcal{D})$ with a value 3 and \mathcal{SHIF} with value 6—see supplementary data). With the calculated IMC (recall Eq. 3), the correlation between DL fragment and the IMC is -0.18 for the ontologies in Set1 and -0.74 for the ontologies in Set2. This possibly may change a little by tweaking the values assigned to the DL fragments, but not such as to obtain a strong, meaningful correlation between detected DL fragment and the IMC.

Correlations for several measures are included in Table 2. The only substantial correlations found are between all pitfalls per ontology elements and size (in boldface), although with all pitfalls minus P8, there is no obvious correlation anymore. p-values were computed with the 1-tailed unpaired Student t-test,




Figure 1: Percentage of occurrence of a pitfall in the three sets of ontologies.




Figure 2: Average number of pitfall/ontology, by set; for P8, the averages are 62, 297, and 303, respectively.

which are also included in Table 2. Using a generous $p < 0.05$ for no difference between the number of pitfalls per ontology and DL fragment, IMC, OE, or OS, then the hypotheses have to be rejected mainly for novices (boldface in Table 2). Correlations were also computed for certain pitfalls and relevant ontology elements, as shown in Table 3; e.g., P5 is about inverse relationships, hence, one might conjecture it is correlated with the amount of object properties in the ontology. This only holds strongly for P8 and the elements in the Set2 ontologies, which explains why there are significant correlations for all pitfalls but not all minus P8 in Table 2. A weakly possibly interesting correlation exists for P5, P26, P27 in the Set1 ontologies, and for P13 in the well-known ontologies.

Comparing pitfalls among Set1, Set2, and Set3 with the 1-tailed unpaired Student t-test, then the null hypothesis—no difference—has to be rejected for novice vs. mature if one ignores pitfall P8 ($p = 0.0096$), i.e., one can observe a difference, but this does not hold anymore for all pitfalls ($p = 0.13$). The results are inconclusive for the other combinations: all pitfalls novice vs. random $p = 0.15$, all mature

vs. random $p = 0.98$, all minus P8 novice vs. random $p = 0.37$, and all minus P8 mature vs. random $p = 0.82$.

3.2.2 Qualitative analysis of the detected pitfalls

As the pitfalls in the catalogue (and thus OOPS!) are relatively coarse-grained, we examined the OOPS!-detected pitfalls of the ontologies in Set1 and Set2 on the correctness of detection. That is, although the algorithms in OOPS! are implemented correctly, they may detect more pitfalls than what an ontology developer may see as a problem, and such insights may, on the one hand, help refining a pitfall and, on the other hand, downgrade a pitfall to being irrelevant practically. Of the analyses carried out (included in the supplementary data file), we highlight four types of pitfalls that illustrate well an aspect of ontology development practices (P4), subject domain peculiarities (P7), language features (P13), and modelling (P24).

P4: unconnected ontology elements. OOPS! correctly identifies ‘orphan’ classes and properties, but they are debatable in some cases. For instance, an or-

Table 1: Totals for the three sets of ontologies (rounded off), with and without the annotation pitfall (P8).

Ontology Pitfalls	Set1: Novices		Set3: Random		Set2: Well-known		Combined	
	All	All – P8	All	All – P8	All	All – P8	All	All – P8
Total	2046	626	133746	26330	7639	1277	143436	28238
Minimum	23	3	0	0	15	2	0	0
Maximum	366	95	7948	1999	3920	207	7948	1999
Average	89	27	735	145	364	61	353	70
Median	65	19	50	14	137	48	54	16
St. dev.	74	26	1147	244	846	53	1101	231

Table 2: Correlations and p-values for specific pitfalls and ontology size and complexity, with the relatively interesting values in boldface; p/o = pitfalls/ontology; DL = DL fragment; where $p < 0.0001$, only 0 is written in the cell due to width limitations.

Set p/o	Set1: Novices				Set2: Well-known				Both			
	All		All – P8		All		All – P8		All		All – P8	
	Corr.	p	Corr.	p	Corr.	p	Corr.	p	Corr.	p	Corr.	p
DL	0.33	0	0.18	0.0002	0.49	0.066	0.52	0	0.38	0.020	0.38	0
IMC	0.06	0	-0.14	0	-0.21	0.056	-0.36	0	-0.14	0.017	-0.2	0
OE	0.998	0.47	0.70	0.0003	0.993	0.84	0.57	0.068	0.990	0.79	0.58	0.025
OS	0.58	0.0072	0.67	0	0.998	0.34	0.52	0.10	0.995	0.24	0.52	0.044

Table 3: Correlations by pitfalls and ontology metric, with the most interesting values in boldface, and potentials in italics; “–”: no pitfall detected, hence, no correlation.

Ontology correlation	Set1	Set2	Both
P5 – no. OPs	<i>0.71</i>	0.52	0.58
P11 – no. OPs	0.41	0.34	0.40
P13 – no. OPs	0.54	<i>0.78</i>	<i>0.77</i>
P25 – no. OPs	0.36	0.30	0.32
P26 – no. OPs	<i>0.72</i>	-0.25	0.25
P27 – no. OPs	<i>0.71</i>	0.61	0.59
P28 – no. OPs	–	–	–
P29 – no. OPs	-0.20	0.15	-0.03
P3 – no. C	–	–	–
P6 – no. C	–	–	–
P7 – no. C	0.17	0.04	0.06
P10 – no. C	-0.08	-0.13	-0.09
P21 – no. C	–	-0.10	-0.06
P24 – no. C	0.15	-0.04	0.01
P8 – no. C	0.22	0.9975	0.9909
P8 – no. OE	0.51	0.9899	0.9848

phan’s subclasses are used in a class expression, i.e., the orphan class is used merely as a way of grouping similar things alike a so-called ‘abstract class’ in UML. The Deprecated and Obsolete orphans are typically present in bio-ontologies, which is deemed a feature in that field. A recurring type of orphan class was to add a class directly subsumed by owl:Thing to indicate the subject domain (e.g., a Bakery class for an ontology about bakery things), which might be defensible in a distributed ontology, but not in a sin-

gle domain ontology. Overall, each of these practices require a more substantive argument whether they deserve to be a false positive or not.

P7: Merging different concepts in the same class. OOPS! detects a few occurrences that are false positives, besides the many correctly identified ones. For instance, a RumAndRaisinFlavour of ice cream does not constitute merging different classes, but a composite flavour and would not have been a false positive if that flavour had obtained its own name (e.g., RummyRaisin). From a computational perspective, there is no easy way to detect these differences.

P13: Missing inverse relationships. The issues with inverses are contentious and hard to detect, especially since OWL and OWL 2 differ in their fundamental approach. Unlike OWL, OWL 2 has a feature `ObjectInverseOf`, so that for some object property `hasOP` in an OWL 2 ontology, one does not have to extend the vocabulary with an `OPof` property and declare it as the inverse of `hasOP` with `InverseObjectProperties`, but instead one can use the meaning of `OPof` with the axiom `ObjectInverseOf(hasOP)`. In addition, GFO’s `exists_at` and BioTop’s `abstractlyRelatedTo` do not readily have an inverse name, and a modeller likely will not introduce a new property for the sake of having a named inverse property when it is not needed in a class axiom. Overall, P13 is detected more often than warranted from a modeller’s viewpoint, and it could be refined to only those cases where the declaration of `InverseObjectProperties` is missing; e.g., both `manufacturedBy` and `hasManufacturer` are in the car ontology, but they are not declared inverse though they clearly are, which OOPS! detects already.

P24: Using recursive definition. This pitfall is tricky to define and detect. In general, recursive definitions are wrong, such as the pattern $X \equiv X \sqcap R.Y$, which should be detected, and likewise detecting unintended assertions, such as `CarrotFilling \sqsubseteq \exists hasFillingsAndTopping.CarrotFilling` (in the bakery (novice’s) ontology). However, P24 currently detects whether the class on the left-hand side of the subsumption or equivalence occurs also on the right-hand side, which is not always a problem; e.g., `DM-Process \sqsubseteq \exists has subprocess.DM-Process` in DMOP is fine. These subtle differences are difficult to detect automatically, and require manual inspection before changing or ignoring the pitfall.

Removal of the false positives reduces the observed minor differences between the three sets of ontologies, i.e., roughly equalize the percentages per pitfall. Put differently, this supports the observation that there is a *general landscape of pitfalls*.

3.2.3 Candidate pitfalls

The novices’ ontologies had been analysed manually on modelling mistakes before OOPS! and before consulting the catalogue. In addition to detecting the kind of pitfalls already in the catalogue, new ones were detected, which typically occurred in more than one ontology. We refer to them here as new *candidate pitfalls* (Cs) to add to the catalogue:

- C1. *Including some form of negation in ontology element names.* For example, `DrugAbusePrevention` (discussed in (Schulz et al., 2009)), and `NotAdults` or `ImpossibleHand` (in poker ontology). This pitfall refers to an anomaly in the element naming.
- C2. *Distinguishing between file name and URI.* This is related to naming issues where the `.owl` file has a meaningful name, but the ontology URI has a different name (also observed in (Keet, 2011)).
- C3. *Confusing part-of relation with subclass-of relation.* This pitfall is a special and very common case of pitfall P23 (using incorrectly ontology elements) (see (de Cea et al., 2008)). As part of this pitfall, there is also the case in which the most appropriate part-whole relation in general is not selected (see also (Keet et al., 2012)).
- C4. *Misusing min 1 and some.* This pitfall affects especially ontology feature usage due to the OWL restrictions (note: Protégé 4.x already includes a feature to change all such instances).
- C5. *Embedding possibility/modality in the ontology element’s name.* This pitfall refers to encapsulating a modality (“can”, “may”, “should”) in an element’s name (e.g., `canCook`).

3.3 Discussion

Whilst giving valuable insight in the prevalence of pitfalls in existing ontologies, the results obtained falsify hypotheses (i) (except for novice vs. mature when discounting P8), (ii), and (v), partially validate (iv) (for all pitfalls and mature ontologies), and validate (iii), which is not exactly as one may have expected, and it raises several possible interpretations.

First, the set of pitfalls currently implemented in OOPS! is limited and with more and more refined checks, substantial differences may be found. Perhaps this is the case, but it does not negate the fact that it is not the case for the 21 already examined and therefore not likely once extended. In addition, recently, the notion of good and safe object property expressions has been introduced (Keet, 2012), where manual evaluation with a random set of ontologies—including some of the ones in Set2—revealed advanced modelling issues concerning basic and complex object property expressions. This further supports the notion that, for the time being, there is a general landscape compared to salient differences among levels of maturity.

Second, the well-known ontologies are possibly not really mature and exemplary after all (the converse—that the novices’ ontologies in Set1 are ‘as good as the well-known ones’—certainly does not hold), for they are quite close to the ones in Set3; i.e., that some ontology is widely known does not imply it is ‘good’—or, at least: has fewer pitfalls than an—ontology being developed by a novice ontologist. This makes it more difficult to use them in ontology engineering courses, where one would like to point students to ‘good’ or ‘exemplary’ ontologies: if well-known ontologies have those pitfalls, they are more likely to be propagated by the students “because ontology x does it that way”. This attitude was observed among the novices with respect to P11, because the popular Protégé OWL Pizza tutorial (<http://www.co-ode.org>) advises against declaring domain and range of properties (page 37), which may explain why P11 was detected often.

Third, it may be reasonable to argue that ‘maturity’ cannot be characterised by absence of pitfalls at all, but instead is defined by something else. Such a ‘something else’ may include its usefulness for its purpose—or at least meeting the requirements—or, more abstract, the precision and coverage as introduced by Guarino (see also Fig 2 in (Guarino et al., 2009)). Concerning the latter, this means both a high precision and maximum coverage of the subject domain one aims to represent in the ontology. It is known one can improve on one’s low precision—i.e., the ontology admits more models than it should—by using a more expressive language and adding more

class expressions, but this is easier said than done (except for the new tool to add more disjointness axioms (Ferré and Rudolph, 2012)). For domain ontologies, another option that influences to notion of being well-known and mature is its linking to a foundational ontology and that therewith less modelling issues occur (Keet, 2011), but this has to do with the knowledge that is represented, not with, e.g., language feature misunderstandings. We leave a more detailed investigation in this direction for future works.

4 CONCLUSIONS

We performed a quantitative analysis of the pitfalls developers included in ontologies by analyzing different sets of data obtained after using OOPS!. All implemented types of pitfalls have been detected in the ontologies scanned with OOPS!, but the most common ones are lack of annotations, absence of domain and range axioms, and issues with inverses, and to a lesser extent creating unconnected ontology elements and using a recursive definition. Five new candidate pitfalls have been identified upon closer inspection of the novices' ontologies. Analysis showed that there is no clear evidence of noteworthy differences between ontologies developed by novices, well-known ones, and the random set of ontologies, except for novice vs. mature when disregarding pitfall P8, and for novices, the pitfalls per ontology is related to the size of the ontology complexity of the ontology. Thus, the analysis provides a data-driven general landscape of pitfalls in current ontologies.

We are extending the pitfall catalogue, and are working on a better characterization of 'maturity' in ontologies and how such a characterization is related to the set of most common pitfalls.

ACKNOWLEDGEMENTS

This work has been partially supported by the Spanish projects BabelData (TIN2010-17550) and BuscaMedia (CENT 2009-1026).

REFERENCES

- Brank, J., Grobelnik, M., and Mladenic, D. (2005). A survey of ontology evaluation techniques. In *Proc. of SiKDD 2005*. Ljubljana, Slovenia. 2005.
- de Cea, G. A., Gómez-Pérez, A., Montiel-Ponsoda, E., and Suárez-Figueroa, M. C. (2008). Natural language-based approach for helping in the reuse of ontology design patterns. In *Proc. of EKAW'08*, volume 5268 of *LNCS*, pages 32–47. Springer.
- Ferré, S. and Rudolph, S. (2012). *Advocatus diaboli - exploratory enrichment of ontologies with negative constraints*. In *Proc. of EKAW'12*, volume 7603 of *LNAI*, pages 42–56. Springer.
- Gangemi, A., Catenacci, C., Ciaramita, M., and Lehmann, J. (2006). Modelling ontology evaluation and validation. In *Proc. of ESWC'06*, volume 4011 of *LNCS*, pages 140–154. Springer.
- Gómez-Pérez, A. (2004). Ontology evaluation. In Staab, S. and Studer, R., editors, *Handbook on Ontologies*, pages 251–274. Springer.
- Guarino, N., Oberle, D., and Staab, S. (2009). What is an ontology? In Staab, S. and Studer, R., editors, *Handbook on Ontologies*, chapter 1, pages 1–17. Springer.
- Guarino, N. and Welty, C. (2009). An overview of OntoClean. In Staab, S. and Studer, R., editors, *Handbook on Ontologies*, pages 201–220. Springer, 2 edition.
- Keet, C. M. (2011). The use of foundational ontologies in ontology development: an empirical assessment. In *Proc. of ESWC'11*, volume 6643 of *LNCS*, pages 321–335. Springer.
- Keet, C. M. (2012). Detecting and revising flaws in OWL object property expressions. In *Proc. of EKAW'12*, volume 7603 of *LNAI*, pages 252–266. Springer.
- Keet, C. M., Fernández-Reyes, F. C., and Morales-González, A. (2012). Representing mereotopological relations in OWL ontologies with ONTOPARTS. In *Proc. of ESWC'12*, volume 7295 of *LNCS*, pages 240–254. Springer.
- Noy, N. and McGuinness, D. (2001). *Ontology development 101: A guide to creating your first ontology*. TR KSL-01-05, Stanford Knowledge Systems Laboratory.
- Poveda, M., Suárez-Figueroa, M. C., and Gómez-Pérez, A. (2010). Common pitfalls in ontology development. In *Current Topics in Artificial Intelligence, CAEPIA 2009 Selected Papers*, volume 5988 of *LNAI*, pages 91–100. Springer.
- Poveda-Villalón, M., Suárez-Figueroa, M. C., and Gómez-Pérez, A. (2010). A double classification of common pitfalls in ontologies. In *Proc. of Workshop on Ontology Quality*, CEUR-WS. Co-located with EKAW'10.
- Poveda-Villalón, M., Suárez-Figueroa, M. C., and Gómez-Pérez, A. (2012). Validating ontologies with OOPS! In *Proc. of EKAW'12*, volume 7603 of *LNAI*, pages 267–281. Springer.
- Rector, A. et al. (2004). OWL pizzas: Practical experience of teaching OWL-DL: Common errors & common patterns. In *Proc. of EKAW'04*, volume 3257 of *LNCS*, pages 63–81. Springer.
- Roussey, C., Corcho, O., and Vilches-Blázquez, L. (2009). A catalogue of OWL ontology antipatterns. In *Proc. of K-CAP'09*, pages 205–206.
- Sabou, M. and Fernandez, M. (2012). Ontology (network) evaluation. In *Ontology Engineering in a Networked World*, pages 193–212. Springer.
- Schulz, S., Stenzhorn, H., Boekers, M., and Smith, B. (2009). Strengths and limitations of formal ontologies in the biomedical domain. *Electr. J. Communication, Information and Innovation in Health*, 3(1):31–45.
- Vrandečić, D. (2009). Ontology evaluation. In Staab, S. and Studer, R., editors, *Handbook on Ontologies*, pages 293–313. Springer, 2nd edition.