

Editorial

Proposed guidelines for publishing ontology papers

Roberta Ferrario^{a,*} and Michael Grüninger^b

^a *Institute of Cognitive Sciences and Technologies – CNR – Laboratory for Applied Ontology, Trento, Italy*

^b *Department of Mechanical and Industrial Engineering, University of Toronto, Ontario, Canada M5S 3G8*

1. Introduction

The discipline of Applied Ontology is facing several challenges. Some of these challenges concern ontologies themselves, for instance the ability to appropriately and exhaustively represent certain domains, or to correctly and efficiently perform reasoning; others are more tied to their use, for example how they are used in applications and with which benefits, or whether and how it is possible to align them with other existing ontologies. What seems evident is that, in order to be widely and profitably used, they should demonstrate a high quality level and this is unfortunately not always the case.

All of this is reflected in our experience with the Journal. Many papers are submitted on domain-specific ontologies, but few are accepted because they are not meeting quality standards. Very rarely the ontologies they present are aligned or even inspired by upper level ontologies and, worse, quite often the presentation of the ontology consists solely of a collection of Protégé screenshots. The ontology has a poor axiomatization, and even if axioms are presented, we find mere translation of them into natural language, without justification or explanation. The papers contain at best a shallow analysis of the ontology, without a substantial foundation in underlying theories or formal techniques and without a sound motivation of the adopted modeling choices. Finally, there is no comparison with similar/alternative approaches in the same domain of application. As a consequence, any possibility of exchanging information with systems based on different ontologies of the same domain is precluded from the start.

More broadly, these problems draw attention to general questions the community should try and address. How do we measure progress in Applied Ontology? What is considered an advancement in the field? What are we doing as a community, and what role does the Journal play? We would like this editorial to initiate such a discussion.

*Corresponding author. E-mail: roberta.ferrario@loa.istc.cnr.it.

2. Crisis of content

Vociferous battles are waged over upper ontologies, yet domain ontologies are often neglected within conferences such as FOIS. On the other hand, most users of ontologies are concerned with quite specific domains of application, such as biology, medicine, engineering, the Internet of Things. These disciplines look to the Applied Ontology community for assistance in the design and analysis of ontologies. Can the Journal and the community through it play a role in providing this assistance?

The Crisis of Content that we face arises from a tragic irony – ontologies are supposed to be sharable and reusable, but is this really happening in practice? For example, if someone wants to use an ontology of everyday activities, how do they know whether such an ontology even exists? How can people find ontologies that they are searching for? On the other hand, in some domains there is a plethora of ontologies with no understanding of their relationships – are we continually reinventing the wheel? If someone finds an ontology that is potentially relevant, can it be combined with other ontologies?

2.1. FAIR ontologies

In 2016, the FAIR principles were first articulated by Wilkinson et al. (2016) with the aim of favoring the use, management and dissemination of scientific data, metadata and infrastructures. Such principles prescribe that data, metadata and infrastructure should be:

- *Findable: easy to identify and find for both humans and computers, with metadata that facilitate searching for specific datasets,*
- *Accessible: stored for long term so that they can easily be accessed and/or downloaded with well-defined access conditions, whether at the level of metadata, or at the level of the actual data,*
- *Interoperable: ready to be combined with other datasets by humans or computers, without ambiguities in the meanings of terms and values,*
- *Reusable: ready to be used for future research and to be further processed using computational methods. This requires adequate information about how the data were obtained and processed (provenance) and an appropriate license.*

To address the Crisis of Content, we propose that the Applied Ontology community adapt these principles to the ontologies that appear in papers in the Journal.

2.1.1. Accessible ontologies

Ontologies are often scattered across the Semantic Web. In order to make ontologies accessible, we recommend that all ontologies that are referenced in a paper in the Journal should be available in an ontology repository on the Web, such as BioPortal,¹ OntoHub,² or COLORE.³ In the Draft Report “A 20-Year Community Roadmap for Artificial Intelligence Research in the U.S.”, the American Association for Artificial Intelligence has already identified the critical role to be played by open knowledge repositories.

2.1.2. Findable ontologies

If we are to easily identify and find ontologies, ontology repositories alone will be inadequate if there is insufficient support for metadata of the ontologies within the repositories. We invite the hosts of

¹<https://bioportal.bioontology.org/>

²ontohub.org

³colore.oor.net

the existing ontology repositories to adopt standard metadata to facilitate the search for ontologies and explore new ways of organizing repositories to make it easier for humans to find ontologies in similar domains. Ideally, metadata about ontologies should remain accessible even when ontologies are for some reason no longer in the repository.

2.1.3. Interoperable ontologies

The axiomatization of an ontology in a standard language such as OWL or Common Logic together with open implementations of translators between these languages has gone a long way to making ontologies interoperable, the adoption of such languages is thus encouraged. Similarly, whenever possible, the alignment to upper ontologies is a practice that greatly enhances semantic interoperability and constitutes a desideratum.

Ontologies may also be used to make data interoperable, by providing a semantic infrastructure to datasets, so that even heterogeneous data deriving from different datasets can be integrated on the basis of the meaning ontologies ascribe to their content. Just as we have recommended that ontologies be available, the datasets used with ontologies in their applications should also be made available within repositories. This would support more effective comparison of alternative ontologies that have overlapping conceptual coverage.

2.1.4. Reusable ontologies

To enable an effective reuse of ontologies, researchers should be in the condition of finding already existing ontologies of the same domain either to reuse them directly, or to reuse some of their modules and extend or modify them. As already mentioned, if ontologies are archived in well structured and indexed repositories, they should already be findable and accessible. These are necessary requirements that are presupposed for reuse, but they are not sufficient. What is further needed is that the ontologies contained in repositories are curated, with standard metadata that can help researchers in comparing and choosing the domain ontology that best marches their requirements and can thus be reused.

3. The crisis of quality

Ontologies are more than software artifacts – they can be formally evaluated and compared. There is currently no consensus within the community about the evaluation criteria for ontologies, particularly for more domain-specific ontologies. As a result, there is wide variability among the different ontologies that have been proposed. This problem is compounded by the lack of clear criteria for presenting ontologies within publications.

What makes Applied Ontology unique is its interdisciplinary nature, combining philosophy, logic, linguistics, and computer science. Each of these disciplines contributes its own set of criteria for evaluating ontologies, and any coherent approach to ontology evaluation used in reviewing a paper needs to take these perspectives into account. Furthermore, each domain of application implicitly poses a different validation problem for domain ontologies, independently of the philosophical and logical techniques for ontology verification.

3.1. Publication of ontology papers

In addition to adopting FAIR Principles for ontologies, we are also proposing guidelines for evaluating papers that present specific ontologies. These guidelines are meant to help editors and reviewers in

their evaluation work but, most importantly, also authors who are not very familiar with the Journal in understanding what is expected from their submissions.⁴

Papers should contain an adequate description of:

Requirements. At Ontology Summit 2013 Neuhaus et al. (2013) proposed an ontology lifecycle that encompasses requirements development, ontological analysis, design, and deployment. Requirements are typically presented as competency questions, but this can also include use cases that present scenarios describing the problem being solved and motivation for why the ontology has been designed.

Axiomatization (ontology content). If a paper that presents an ontology does not explicitly contain the full axiomatization of the ontology, the community is not able to properly understand and evaluate the claims made by the authors.

Ontological choices. Axiomatization necessarily implies ontological choices. It is important that choices are explicitly stated, so that the reader can evaluate whether they subscribe to the choice and then can adopt the ontology for their purpose and verify if the axioms express correctly the declared choice.

Implementation. Irrespective of the application of an ontology, there needs to be documentation of how the ontology is being used. For example, is a reasoner being used to infer conclusions from the axioms? Are existing algorithms being augmented with the use of an ontology? Is the ontology being used to support the semantic integration of datasets?

Evaluation. Following the results of Ontology Summit 2013, we can consider both intrinsic and extrinsic criteria for ontology evaluation that should be covered in papers about ontologies. Intrinsic evaluation focuses on the axiomatization of the ontology alone, including consistency. Extrinsic evaluation considers not only the ontology, but also the applications of the ontology, such as semantic integration, decision support, and search.

Rigorous comparison with other related ontologies. In many cases, an ontology presented in a Journal paper is not the first to be proposed for a given domain. Best practice in academic writing includes a review of existing work, but this needs to be carried further. Formal logical and ontological analysis of the relationship to existing ontologies needs to be included. Implicit in such an analysis is a response to the question of why a new ontology is needed, and includes a summary of the expected and realized benefits of using the proposed ontology.

4. Looking forward

A first step we take as editors-in-chief of the Journal is a renovation of the Editorial Board, with new additions of scholars who will bring their outstanding expertise in topics as ontology of medicine, classical physics, mental and social properties, manufacturing, digital humanities, cultural heritage; in ontology learning, mapping, aligning, integration, aggregation, evaluation, reuse and many more. We take the opportunity with this editorial to thank them for having accepted to be part of the board and to warmly welcome them all; at the same time, we would also like to gratefully thank the resigning members for their continuous and fruitful collaboration and support through the years.

⁴There are, of course, many papers in the Journal which are not about specific ontologies, but rather are concerned with broader issues of philosophical and ontological analysis. The guidelines being proposed in this Editorial do not necessarily apply to them.

References

- Neuhaus, F., Vizedom, A., Baclawski, K., Bennett, M., Dean, M., Denny, M., Gruninger, M., Hashemi, A., Longstreth, T., Obrst, L., Ray, S., Sriram, R., Schneider, T., Vegeffi, M., West, M. & Yim, P. (2013). Towards ontology evaluation across the life cycle. *Applied Ontology*, 8(3), 179–194. doi:[10.3233/AO-130125](https://doi.org/10.3233/AO-130125).
- Wilkinson, M., Dumontier, M., Aalbersberg, J. & Appleton, G. (2016). The FAIR guiding principles for scientific data management and stewardship. *Scientific Data*, 3, 160018. doi:[10.1038/sdata.2016.18](https://doi.org/10.1038/sdata.2016.18).