



<http://www.diva-portal.org>

Postprint

This is the accepted version of a chapter published in *Ontology Engineering with Ontology Design Patterns*.

Citation for the original published chapter:

Hammar, K., Blomqvist, E., Carral, D., Van Erp, M., Fokkens, A. et al. (2016)

Collected Research Questions Concerning Ontology Design Patterns.

In: Pascal Hitzler, Aldo Gangemi, Krzysztof Janowicz, Adila Krisnadhi, Valentina Presutti (ed.),

Ontology Engineering with Ontology Design Patterns (pp. 189-198). IOS Press

Studies on the Semantic Web

<http://dx.doi.org/10.3233/978-1-61499-676-7-189>

N.B. When citing this work, cite the original published chapter.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:hj:diva-32135>

Collected Research Questions

Karl Hammar¹, Eva Blomqvist², David Carral³, Marieke van Erp⁴,
Antske Fokkens⁵, Aldo Gangemi⁶, Willem Robert van Hage⁷, Pascal
Hitzler⁸, Krzysztof Janowicz⁹, Nazifa Karima¹⁰, Adila Krisnadhi¹¹,
Tom Narock¹², Roxane Segers¹³, Monika Solanki¹⁴, and Vojtech
Svatek¹⁵

¹*Jönköping University, Linköping University*

²*Linköping University*

³*Wright State University*

⁴*Vrije Universiteit Amsterdam*

⁵*Vrije Universiteit Amsterdam*

⁶*CNR-ISTC*

⁷*Vrije Universiteit Amsterdam*

⁸*Wright State University*

⁹*University of California, Santa Barbara*

¹⁰*Wright State University*

¹¹*Wright State University*

¹²*Marymount University*

¹³*Vrije Universiteit Amsterdam*

¹⁴*University of Oxford*

¹⁵*University of Economics, Prague*

This chapter lists and discusses open challenges for the ODP community in the coming years, both in terms of research questions that will need to be answered, and in terms of tooling and infrastructure that will need to be developed to increase adoption of ODPs in academia and industry. The chapter is organised into three sections: Section 1 focuses on issues pertaining to the patterns themselves, including understanding their features and qualities, and developing pattern languages and standards. Section 2 concerns the evaluation and development of methods for using, constructing, and extracting ODPs. Finally, Section 3 focuses on tooling and infrastructure development, including Ontology Engineering environments that support pattern use, pattern repository development, and sustainability and versioning issues.

1 Patterns and Standards

1.1 Quality and Evaluation

Many ontologies are rarely, if ever, reused by third parties. This flies in the face of one of the original main motivations for creating ontologies, namely as shared conceptualizations. However, the exact reasons why ontologies see so little reuse are at this point very poorly understood. Anecdotal evidence indicates that ontologies often are poorly designed and constructed, difficult to understand, insufficiently documented and maintained, too specific or too broad (or both). Social dynamics and findability may also play a role. As ontology design patterns are created and used, it behooves the community to really understand the causes for the lack of reuse.

Without quality documentation, it is nearly impossible for ontology engineers to discover existing reusable ODPs, whether by browsing manually or using some search tool. Thus understanding how to best document ODPs (which documentation components are needed, in which order should they be displayed, using which graphical layout, etc) and developing benchmarks for evaluating ODP documentation are important research directions.

Further, the ODP community needs to address possible causes related to poor design, to understandability issues, and to finding an appropriate balance for specificity versus generality. A thorough and evidence-supported understanding of these issues, e.g. what exactly poor design and good design are, appears to be necessary to advance on these fronts; investigations into ontological anti-patterns fall into this realm [15]. In particular, there seems to be a lack of user-centric evaluations addressing the ontology reuse issue, with only a few notable exceptions, e.g. [7].

In addition to the above, even if a greater understanding of ODP design and documentation are developed, we still need to challenge and evaluate our basic assumptions. We often take as a given that properly supported, ODP use will foster reuse of ontologies, help clean up messy data and facilitate data integration, when in fact, these claims are to a large degree still unverified hypotheses. An important step forward for the ODP community would be to show that the veracity of such claims. This could be done via empirical studies such as [7] and the development of useful applications which make use of the ODP paradigm.

1.2 Pattern Languages and Standards

Currently, ontology design patterns are mostly presented in the form of OWL files, with some minimalistic accompanying documentation. While OWL can indeed be used to express the core of a pattern, it does not provide native ways to represent additional information which would be helpful for reuse or organization of patterns.

For example, it would seem to be important to understand in depth (and be able to represent formally) how different patterns relate to each other. This is important to understand compositionality aspects, e.g. how to create an ontology based on existing ontology design patterns. Key relations are specialization, generalization, and composition and are discussed as in [13], but often relationships are more complex. For example, there are notions of views or shortcuts, see [6], which can be understood as a type of temporary simplification of a pattern.

Using OWL to express axiomatizations provided with patterns is also limiting because the open-world assumption underlying OWL does not cater for expressing non-monotonic constructs like constraints, and because some patterns call for an axiomatization of, say, transitive closure or general rules not expressible in OWL DL. Some first steps towards developing such a representation language have already been undertaken [4] but more work remains to be done.

Another challenge concerns naming in ODPS – currently, the lexical aspect of patterns is rather ad hoc. There has already been some work specifically devoted to ontological naming conventions [16] and even mentioning cross-entity structures [18], but little long-term systematic work has been carried out, and the respective category of the ODP portal is, consequently, empty¹. The naming aspect is important both for content patterns (which are a kind of mini-ontologies with their own vocabulary) and for logical patterns, where structural transformation from one modeling variant (e.g., reification or meta-modeling) may also involve transformation of naming such as generation of derived word forms of different parts of speech.

2 Method Development

2.1 Usage Methods

Currently there is only one ontology engineering methodology that explicitly mentions the use of ODPs, i.e. the eXtreme Design (XD) methodology [11, 10]. If we compare this situation that in software engineering we find that in that discipline, almost regardless of what methodology you use, you will take design patterns into account in some way or another. Following this example, in ontology engineering it would probably be beneficial to be able to use ODPs in any ontology engineering methodology, and not only XD. This raises the questions of how ODP usage would look like in other kinds of ontology engineering methodologies, and if different variations of the ODP concept may be needed there?

We also note that there are still open questions relating to the improvement and further development of the XD methodology. Improvement sug-

¹<http://ontologydesignpatterns.org/wiki/Category:NamingOP>

gestions include providing ways of tailoring the methodology to different project settings such as the staffing situation, duration of the project, the overall goal of the project (e.g. type of ontology, how much focus is put on correct concept definitions and terminology vs. logical structure and reasoning capabilities), and also detailing existing steps of the methodology to better fit realistic project settings. The latter includes the question of reuse - how do you reuse an existing knowledge structure (potentially an existing ontology) and extend that by using ODPs? How do you know an ODP fits your existing structure, so that you know what ODPs to select for your extensions?

2.2 ODP Development

The advance of ontology design patterns is caught in a catch-22: In order to provide convincing evidence for the added value of ODPs, the community requires access to a well-organized, well-documented, and well-maintained set of interlinked high-quality ontology design patterns. At the same time, however, there is a lack of incentive (and funding) to provide these, as long as this added value has not yet been convincingly demonstrated.

A joint effort to create, document, and properly catalogue key ontology design patterns will be needed, as well as a discussion on which patterns are needed and how they should be provided. A pattern creation effort is ongoing for the geosciences domain, in form of the U.S. GeoVocamps [6], but the process is still very ad-hoc, and a broader organized effort will be needed.

Good, and highly visible, use cases can be powerful drivers of such community efforts. While use cases for pattern-driven ontology engineering exist (e.g. the recent [8, 9]), they need to be catalogued, made accessible, and assessed regarding the relevance of ODPs to the scenario. Potential additional use cases could be developed with focus on data publishing and reuse, e.g. in the realm of linked data [14], which has seen much work in recent years. A regular challenge or competition event at the WOP workshops could be set up to facilitate this development.

One recent area where the development of more ODPs could fill a clear need would be when applying some non-standard features in ontologies. There are many tutorials, courses, modelling guidelines and ODPs that can help when using standard OWL. However, every year new extensions to the standards are developed, e.g. in terms of managing and reasoning over streaming data, in terms of modelling and reasoning over uncertainty etc. Nevertheless, there are so far few ODPs for such extensions to standard OWL, although this is an area where help and support for the ontology engineer would be really badly needed.

Another challenge concerns methods and best practices for the development of new patterns, which is currently undertaken in a rather ad-hoc

manner. Just as guidelines are available for publishing ontologies and linked data, best practices must be developed that guide pattern developers in a systematic way to address the problem of developing a design pattern. These guidelines must empower the pattern developers to ascertain when their pattern axiomatisation is indeed a pattern and when it moves into the realm of a full fledged ontology before they publish it. Best practices may include defining the scope of the pattern, the minimum number of use cases that need to be made available as examples of the pattern application, the level and kind of axiomatisation that is appropriate and the possible extension points for future exploitation of the pattern amongst others, etc.. The inclusion of patterns in the patterns repository should be governed by validating these criteria.

In addition to traditional OWL-based ODPs, we should consider that ontology patterns (especially content/knowledge patterns) exist in more or less explicit forms in datasets and models developed within diverse disciplines beyond ontology engineering and the Semantic Web, including formal and computational linguistics, lexical resources, data modelling, stylesheets, microformats, web pages, HCI, tables, etc. Making these implicit patterns explicit, and finding links is a big challenge, which we have just started. Cf. [3] for a related manifesto. In particular, the patterns emerging from NLP and linguistic resources, and those from ontology engineering and the Semantic Web have important relations, in terms of grounding, empirical evidence, cognitive relevance, etc. (cf. the NLP formal integration work enabled by FRED [12], and the correspondences found in FrameBase).

2.3 ODP Extraction

The development of Software Engineering design patterns has conventionally followed a bottom-up approach where recurring patterns are extracted from existing software. This also serves as a validation of the usage of the pattern. However, the development of ODPs has largely happened in a top-down manner. This was justified in the early days of ontology development due to a lack of critical mass. However, strategies are now required that encourage bottom-up development and enable automated extraction of patterns from a corpus of ontologies.

Such a corpus, that is as of yet relatively untapped, in fact exists: many well known and widely used ontologies were developed before ODPs were introduced. Some of these ontologies are upper level such as BFO and SUMO, while others such as SNOMED-CT and CIDOC-CRM are relatively large. These ontologies potentially encode a number of ODPs, a fact which is implicitly validated by their widespread usage. Analyzing and restructuring these ontologies to identify patterns using traditional manual methods is a non-trivial task requiring significant support, which is often not available. The development of automated extraction strategies and tools could however

uncover many currently hidden ODPs, thereby bridging the gap between the demand and availability of ODPs, and overcoming the catch-22 discussed in Section 2.2.

3 Tooling and Infrastructure

3.1 Pattern-Capable Ontology IDE:s

Which kind of tool support that is needed and best suited for ODP development and use, is an open question. Some tools have been proposed (such as plugins for the NeOn toolkit, and more recently for WebProtégé [5]) but they do not cover all possible tasks where ODPs can be applied, and in particular there are currently no tools supporting the construction or extraction of patterns, and very few that support sharing, discussing and collecting them (while some online portals do exist, there are no actual client-side tools).

Also, the tooling that does exist is very traditional, introducing the ODPs as a small variation of the normal workflow. An interesting question is whether there are any opportunities for more disruptive changes in tooling that could be supported by ODPs. For instance, (content) ODPs are sometimes viewed as small puzzle pieces, but so far there has not been any attempts to create tools that truly operate on the level of ODPs as their primitives (rather than OWL languages primitives). One could imagine taking ontology engineering to an entirely new level, where the underlying logical language is hidden behind a pattern language. (See also the discussion on pattern language in Section 1.2)

Additionally it is important for the modeler to become aware of the availability of a suitable pattern when trying to model a certain state of affairs – currently there is lack of pattern discovery services in general, as observed in [1] as one of the ODP adoption barriers. While keyword search may be of some use for content patterns it is insufficient for logical-structural patterns. The proposed development of expressive pattern representation language is important but does not suffice by itself if modeling situations are not described as well. Ideally, the retrieval of pattern should be closely interconnected with the support of its reuse. When the modeling challenge mainly consists in expressivity constraints of the target language (e.g., related to n-ary relationships or meta-typing), the modeled situation can possibly be described using first-order logic or a dedicated language such as PURO [17], and starting from this representation, the reuse of the pattern might be invoked semi-automatically for the target ontology [2]. If the problem is of more complex nature, some kind of interactive navigation over some ontology of ontological modeling problems to a node referencing adequate patterns might suit better, and the reuse itself is then likely to be more manual as well. In any case, such supportive technologies are currently either missing or immature.

3.2 Repository Development

Ontology patterns can be needed in many different lexical and structural contexts. An advanced semantic search capability would boost their usage, but e.g. linking competency questions to potentially useful patterns is not trivial, especially if the structure of the question has to be taken into account. There is room here to join in with research on question-answering on linked data.

3.3 Sustainability Challenges

The provision, maintenance and documentation of ontology design patterns are currently done in a very ad-hoc manner. While the community has a rudimentary portal², it lends only limited support and structure. Versioning aspects and other software engineering related issues need to be addressed, and better tool support for pattern-driven ontology engineering needs to be developed. Sustainability related research questions are about finding out how to best implement infrastructure and apply structured and qualitative software engineering principles to ODPs. What would a tool for the efficient support of ODP-driven ontology-engineering look like? How should documentation be provided? What does versioning mean for ontologies or ODPs on the Web?

References

- [1] E. Blomqvist, P. Hitzler, K. Janowicz, A. Krisnadhi, T. Narock, and M. Solanki. Considerations regarding ontology design patterns. *Semantic Web*, 7(1):1–7, 2015.
- [2] M. Dudáš, T. Hanzal, V. Svátek, and O. Zamazal. Obowlmorph: Starting ontology development from puro background models. To be published via Springer LNCS.
- [3] A. Gangemi and V. Presutti. Towards a pattern science for the semantic web. *Semantic Web*, 1(1, 2):61–68, 2010.
- [4] G. Guizzardi. Ontological patterns, anti-patterns and pattern languages for next-generation conceptual modeling. In *Conceptual Modeling*, pages 13–27. Springer, 2014.
- [5] K. Hammar. Ontology design patterns in webprotégé. In *14th International Semantic Web Conference (ISWC-2015)*. CEUR Workshop Proceedings, 2015.

²<http://ontologydesignpatterns.org>

- [6] P. Hitzler, K. Janowicz, and A. Krisnadhi. Ontology modeling with domain experts: The geovocamp experience. In *Proceedings of the Diversity++ Workshop co-located with the 14th International Semantic Web Conference, ISWC*, 2015.
- [7] C. M. Keet. The use of foundational ontologies in ontology development: an empirical assessment. In *The Semantic Web: Research and Applications*, pages 321–335. Springer, 2011.
- [8] A. Krisnadhi, Y. Hu, K. Janowicz, P. Hitzler, R. Arko, S. Carbotte, C. Chandler, M. Cheatham, D. Fils, T. Finin, et al. The geolink modular oceanography ontology. In *The Semantic Web-ISWC 2015*, pages 301–309. Springer, 2015.
- [9] A. A. Krisnadhi, K. J. Yingjie Hu, P. Hitzler, R. Arko, S. Carbotte, C. Chandler, M. Cheatham, D. Fils, T. Finin, P. Ji, M. Jones, N. Karima, K. Lehnert, A. Mickle, T. Narock, M. O’Brien, L. Raymond, A. Shepherd, M. Schildhauer, and P. Wiebe. The geolink framework for pattern-based linked data integration. In *Proceedings of the ISWC 2015 Posters & Demonstrations Track co-located with the 14th International Semantic Web Conference (ISWC-2015)*, 2015.
- [10] V. Presutti, E. Blomqvist, E. Daga, and A. Gangemi. Pattern-based ontology design. In *Ontology Engineering in a Networked World*, pages 35–64. Springer, 2012.
- [11] V. Presutti, E. Daga, A. Gangemi, and E. Blomqvist. extreme design with content ontology design patterns. In *Workshop on Ontology Patterns*, page 83, 2009.
- [12] V. Presutti, F. Draicchio, and A. Gangemi. Knowledge extraction based on discourse representation theory and linguistic frames. In *Proceedings of the 18th International Conference on Knowledge Engineering and Knowledge Management, EKAW’12*, pages 114–129, Berlin, Heidelberg, 2012. Springer-Verlag.
- [13] V. Presutti and A. Gangemi. Content ontology design patterns as practical building blocks for web ontologies. In *Conceptual Modeling-ER 2008*, pages 128–141. Springer, 2008.
- [14] V. Rodriguez-Doncel, A. A. Krisnadhi, P. Hitzler, M. Cheatham, N. Karima, and R. Amini. Pattern-based linked data publication: The linked chess dataset case. In *Proceedings of the 6th International Workshop on Consuming Linked Data co-located with 14th International Semantic Web Conference (ISWC 2105), Bethlehem, Pennsylvania, US*, 2015.

- [15] T. P. Sales and G. Guizzardi. Ontological anti-patterns: Empirically uncovered error-prone structures in ontology-driven conceptual models. *Data & Knowledge Engineering*, 99:72–104, 2015.
- [16] D. Schober, I. Tudose, V. Svátek, and M. Boeker. Ontocheck: verifying ontology naming conventions and metadata completeness in protégé 4. *Journal of biomedical semantics*, 3(2):1, 2012.
- [17] V. Svátek, M. Homola, J. Kluka, and M. Vacura. Mapping structural design patterns in owl to ontological background models. In *Proceedings of the seventh international conference on Knowledge capture*, pages 117–120. ACM, 2013.
- [18] V. Svátek, O. Šváb-Zamazal, and V. Presutti. Ontology naming pattern sauce for (human and computer) gourmets. In *Workshop on Ontology Patterns*, pages 171–178, 2009.