Chapter 4

On the Roles of Logical Axiomatizations for Ontologies

Pascal Hitzler, Data Semantics Laboratory, Wright State University, Dayton, OH, USA

Adila Krisnadhi, Data Semantics Laboratory, Wright State University, Dayton, OH, USA; and Faculty of Computer Science, Universitas Indonesia

In this brief chapter, we will elaborate on the different roles which logical axiomatizations can play for ontology design patterns and for ontologies in general. While doing this, we also encounter some of the many open research questions regarding this issue.

4.1. Logic Choices

Logic-based knowledge representation has a long-standing history [4] and can, e.g., be traced back to Aristotle's *Organon.* As part of Computer Science, it has been a mainstay of the Artificial Intelligence field since its inception in the 1950s.

The resulting breadth and depth of research, results, and applications led to the adoption of logic-based knowledge representation as the foundation for information representation languages on the World Wide Web [6], and thus for the Semantic Web, including the W3C standards RDF [18], RIF [9], and OWL [5].

These standards, of course, capture only the essence of logic-based knowledge representation, which could be (and sometimes is) used for knowledge represenation on the Web. In fact, the discussions are ongoing on suitable formalisms or modifications and extensions of the existing standards [7]. Let us briefly mention some of the key dimensions of these investigations.

Monotonic versus mon-monotonic logics, open versus closed world assumption. First a clarification: Monotonicity (meaning, additional information cannot cause the withdrawal of prevously derived logical consequences) is mostly associated with the open world assumption, while non-monotonicity is mostly associated with the closed world assumption. The distinction is not quite as clear a cut upon a closer look though [20].

Description Logics [17], and thus the Web Ontology Language OWL [5] are based on the *open world assumption* (and are monotonic). Intuitively this means that *absent* information is viewed as *unknown*, e.g., a listing of the eight planets of the solar system does not preclude that there may be any additional ones out there [2]. This choice seems very sensible from a *Web* perspective where additional information may be as easy to obtain as by crawling a few more Web pages.

However, some situations call for an invocation of the *closed world assumption*, the most obvious perhaps being that of database access, where the information in the database is to be considered *complete*. E.g., if a person is not listed as booked on a particular flight by the airline's internal information system, then it should be assumed that the person is in fact *not* booked on this flight. Some of these cases can in fact already be captured by the use of *nominals* as provided by description logics and OWL; and this essentially constitutes a type of world closure captured within a monotonic logic.

More complex cases arise when non-monotonicity is required. Such cases arise, e.g., when the abovementioned database access scenario is coupled with deductive inferences, in the sense that a class shall consist of exactly those instances for which it can be derived deductively that they belong to that class. Assume, for example, that a knowledge base allows us to derive that both David and Raghava are graduate students, but that it does not allow us to derive any information about the graduate student status of Cogan. Under the closed world assumption, we then would this consider a proof that Cogan is in fact *not* a graduate student. If, however, we later add additional information, say that Cogan is a graduate student, then the previous conclusion of Cogan not being a graduate student would have to be retracted. Thus, reasoning under this logic would be non-monotonic.

The question of monotonicity or non-monotonicity of a reasoning task can often be used as a first check whether the task can be expressed in OWL. E.g., in [21] it was investigated whether a releation *more biodiverse than* between regions can be expressed in OWL, if this relative biodiversity is being assessed by means of the the number of different species occurring in these regions. The task is non-monotonic since our assessment may change when further species are found. The example also highlights that non-monotonicity may arise out of some type of *instrospection*, e.g. from reflecting on what is or is not known based on the given data.

While many non-monotonic logics have been defined in the past four decades, three of the initial proposals have arguably been the most prolific: McCarthy's Circumscription [22] which focuses on world closure as key intuition, Moore's Autoepistemic Logic [23] which focuses on introspection, and Reiter's Default Logic [24] which captures reasoning of the type *every bird flies, unless we have evidence for it not to fly.*

In the context of Semantic Web and ontology modeling, non-monotonic issues often arise naturally, as in the biodiversity case or the database case discussed above. While such non-monotonicity cannot be captured by the OWL standard, research investigations are ongoing regarding the possible extension of OWL, or more general description logics, by non-monotonic features. Such extended knowledge representation languages, which combine open and closed world aspects, are often refered to as *local closed world* languages. A significant body of work has been done on such formalisms; rather than repeating earlier enumerations, we simply refer to [10] which contains a relatively recent overview.

Schema-centric versus data-centric approaches. Similar to the above discussed distinction between open and closed world logics, the distinction between schema-centric and data-centric knowledge representation and reasoning approaches is not entirely clear cut. Contrasting them, however still helps in clarifying different emphases of different logics.

On the schema-centric side of the spectrum, e.g., are traditional description lgoics such as $\mathcal{ALC}$ [6] without ABoxes, in which it is not possible to talk about instances (or constants), but which still allow reasoning over class expressions. On the other hand, rule-based approaches such as logic programming [19] and its variants do not natively support reasoning over predicate relationships, but focus on the derivation of facts involving instances (i.e., constants).

This contrast between description logics and, broadly speaking, *rules* has for most of a decade defined a major dividing line (and point of contention) within the Semantic Web field, however more recently this gap is closing [1, 14–16].

Also of relevance in this context is the use of *integrity constraints* in the context of ontologies: While the semantics of OWL axioms is *inferential*, in the sense that it makes it possible to derive (or infer) additional knowledge from the given axioms, integrity constraints constitute constraints which the data has to satisfy in order to be compatible with a schema. E.g., such a constraint could specify that for any person with a bank account a physical address must be on file. Integrity constraints are usually non-monotonic, in the sense that absence of such an address in the example just given would cause an inconsistency, while later addition of an address would reinstate consistency, which cannot happen in a monotonic setting. Mature proposals for adding integrity constraints to OWL have been made [26]. Some aspects of the forthcoming RDF Shapes Constraint Languague (SHACL) [11] may also address integrity constraint issues.

The importance of standardization and tools. A central dividing line between different logics for ontologies on the Web is that some such logics have been standardized, and some have not. Some, such as RDF, RIF and OWL mentioned above, have further been standardized by the World Wide Web Consortium (W3C) for explicit use in a Semantic Web context. As a consequence, such standardized languages have gained significant visibility, and in its wake also a landscape of diverse compatible tools, which range from research prototypes to commercial solutions.

Hence, currently, the Web Ontology Language OWL is the primary language of choice for expressing ontologies. This does not necessarily mean, of course, that it is the best possible conceivable language for this purpose, or that it will not undergo major revisions and extensions in the future. In particular, and as already mentioned above, some things – such as non-monotonic constructs, integrity constraints, some type of rules [13], and other constructs of potential importance for ontology modeling [12] – are not expressible in OWL.

4.2. Axioms for Inferential Reasoning

Arguably, the most obvious use of logical axioms is for inferential reasoning, i.e. to derive logical consequences from information contained in the ontology or knowledge base. Typical examples for schema-centric reasoning would be the derivation of inferred class relationships, e.g., if *Feline* is a subclass of *Mammal*, which in turn is a subclass of *Animal*, then it can be inferred that *Feline* is a subclass of *Animal*. Corresponding data-centric reasoning may infer, e.g., that a *Feline* identified as *Mimi* is also an *Animal*.

The formal semantics of standards such as OWL and RDF in fact is tailored towards inferential reasoning, and significant research efforts are made towards developing ever more efficient algorithms and systems for performing inferential reasoning [3].

Strong inferential reasoning systems for ontologies are of course often tailored towards the standardized representation languages, although exceptions exist, such as the support, by many systems, of the Semantic Web Rule Language SWRL [8], which has only the status of a W3C Member Submission.

4.3. Axioms to Inform Humans

Another, often undervalued use of logical axioms in ontology modeling is that of informing humans about the intended meaning or classes and their relationships. Axioms stated for the purpose of inferential reasoning can of course already help humans to understand the intended meaning of ontology parts, and we therefore would argue that such axioms, in human-readable form, should be an important part of the documentation of an ontology, and not only be published as part of an OWL file.

In the practice of ontology modeling, in fact, we sometimes encounter axioms which quite clearly are not intended for inferrential reasoning, because they constitute tautologies. However, they are informative for humans as to the intended meaning of terms and their relations. A particular case in point would be minimum cardinality statements with the minimum specified as zero, e.g. *A person always has at least 0 children* or *A chess game has at least 0 moves*. Such statements clarify that properties such as *hasChild* or *hasHalfMove* are intended for use together with the classes *Person* and *ChessGame*, respectively, which for example could indicate to a human user that they are not to be used in the context of trees (as data structures) or the game of Go (where it would be more appropriate to use *ply* or *turn*).

Informing humans, and thereby disambiguating the meaning of ontology terms, is also a primary reason why it is sometimes useful to state axioms in documentation which cannot be expressed in the standardized language chosen to formally represent the ontology. This would be the case, e.g. for axioms which can be expressed in first-order predicate logic, but not in description logic fragment which constitutes OWL DL.

And sometimes the axioms governing a part of an ontology can be expressed in OWL, but the result would be rather difficult to read and thus assess by a human. E.g., the example statement *All elephants are bigger than all mice* from [25] can easily be expressed using a (first-order) rule such as

$$\text{Elephant}(x) \wedge \text{Mouse}(y) \rightarrow \text{biggerThan}(x, y),$$

but the corresponding formulation in the description logic corresponding to OWL DL would require the introduction of two new properties, and a total of three axioms:

$$\begin{aligned} \text{Elephant} &\equiv \exists R_1.\text{Self} \\ \text{Mouse} &\equiv \exists R_2.\text{Self} \\ R_1 \circ R_2 &\sqsubseteq \text{biggerThan} \end{aligned}$$

While the OWL version can of course be used in the corresponding OWL file, for the documentation the rule variant may be preferable, since its meaning is more immediately obvious.

4.4. Axioms as Integrity Constraints

An inferential axiom would state that *every person has a birth place*, while a corresponding integrity constraint would state that *every person has a birth place listed* (in this knowledge base). In the first case, if a birth place for a person, say *John*, is not listed or inferable from the data, then this would not constitute an inconsistency: According to the open world assumption, John would still have a birth place, it would just be the case that we don't know about it (yet). In the case of the integrity constraint, absence of a birth place for John would cause an inconsistency (or other exception).

Using OWL (in description logic syntax), the inferential axiom could be expressed, e.g., as

$$\text{Person} \sqsubseteq \exists \text{hasBirthPlace}.\text{Place}.$$

For the integrity constraint version, we know of no standardized ontology language to express it. However, it would be conceivable to, e.g., abuse OWL syntax and use the very same axiom with a disclaimer (for the human reader) that it is to be read as an integrity constraint. This would serve the purpose of informing humans about the intended meaning. The axiom could still be used for inferential semantics of course, it would simply not derive anything if the integrity constraint is satisfied. And a human user interested in checking the integrity constraint would have the option of using other mechanisms for this purpose, e.g., by retrieving all members of the class Person and all triples involving the hasBirthPlace property, and to then check whether each person has a birth place listed.

4.5. Axioms as Shape Constraints

Shape constraints in the context of OWL have not been studied prominently yet. Only recently, they have been taken up by a W3C working group to standardise a shape constraint language for RDF [11]. However, they have often made informal appearances in the form of class diagrams and tools which attempt to create class diagrams from OWL files.

We do not yet sufficiently understand the role of shape constraints for ontologies. Anecdotal experience – from ontologies we have looked at – indicates that domain and range declarations for properties may sometimes be intended as a type of shape constraints, informing the human user about the class diagram in a non-visual form. This is of course in contrast to the standardized semantics, which is inferential also for domain and range declarations. Likewise, tautological axioms such as the minium zero cardinality axioms discussed above, can also play the informal role of shape constraints.

The formalization of shape constraints for OWL ontologies will certainly lead to a non-monotonic formalism, which also indicates that shape constraints cannot be expressed in the montonic Web Ontology Language.

4.6. Conclusion

We have briefly discussed different major approaches to logical knowledge representation as they pertain to theory and practive of ontology modeling and ontology language research. We have also provided four mostly distinct perspectives on the roles which logical axioms can play in order to disambiguate the meaning of ontology constructs or for programmatic use.

We are not aware of any systematic investigations into the different roles axioms do, can, or could play for ontology modeling and practice. However we believe that it should be a worthwhile endeavour to embark on such investigations.

Acknowledgements. This work was partially supported by the National Science Foundation under award 1440202 *EarthCube Building Blocks: Collaborative Proposal: GeoLink – Leveraging Semantics and Linked Data for Data Sharing and Discovery in the Geosciences.*

Bibliography

[1] A. Amarilli and M. Benedikt. Combining existential rules and description logics. In Q. Yang and M. Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 2691–2697. AAAI Press, 2015.

[2] K. Batygin and M. E. Brown. Evidence for a distant giant planet in the Solar system. *The Astronomical Journal*, 151(2):22, 2016.

[3] M. Dumontier, B. Glimm, R. S. Gonçalves, M. Horridge, E. Jiménez-Ruiz, N. Matentzoglu, B. Parsia, G. B. Stamou, and G. Stoilos, editors. *Informal Proceedings of the 4th International Workshop on OWL Reasoner Evaluation (ORE-2015) co-located with the 28th International Workshop on Description Logics (DL 2015), Athens, Greece, June 6, 2015*, volume 1387 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2015.

[4] D. M. Gabbay and J. Woods, editors. *The Handbook of the History of Logic*. Elsevier, Amsterdam, 2004–2014. 11 Volumes as of April 2016.

[5] P. Hitzler, M. Krötzsch, B. Parsia, P. F. Patel-Schneider, and S. Rudolph, editors. *OWL 2 Web Ontology Language: Primer*. W3C Recommendation, 27 October 2009. Available at `http://www.w3.org/TR/owl2-primer/`.

[6] P. Hitzler, M. Krötzsch, and S. Rudolph. *Foundations of Semantic Web Technologies*. Chapman & Hall/CRC, 2010.

[7] P. Hitzler, J. Lehmann, and A. Polleres. Logics for the semantic web. In D. M. Gabbay, J. Woods, and J. Siekmann, editors, *Logic and Computation*, volume 9 of *Handbook of teh History of Logic*, pages 679–710. Elsevier, Amsterdam, 2014.

[8] I. Horrocks, P. F. Patel-Schneider, H. Boley, S. Tabet, B. Grosof, and M. Dean. *SWRL: A Semantic Web Rule Language Combining OWL and RuleML*. W3C Member Submission 21 May 2004, 2004. Available from http://www.w3.org/Submission/SWRL/.

[9] M. Kifer and H. Boley, editors. *RIF Overview*. W3C Working Group Note 22 June 2010, 2010. Available from http://www.w3.org/TR/rif-overview/.

[10] M. Knorr, P. Hitzler, and F. Maier. Reconciling OWL and non-monotonic rules for the Semantic Web. In L. De Raedt, C. Bessiere, D. Dubois, P. Doherty, P. Frasconi, F. Heintz, and P. Lucas, editors, *ECAI 2012, 20th European Conference on Artificial Intelligence, 27-31 August 2012, Montpellier, France*, volume 242 of *Frontiers in Artificial Intelligence and Applications*, pages 474–479. IOS Press, Amsterdam, 2012.

[11] H. Knublauch and D. Kontokostas, editors. *Shapes Constraint Language (SHACL)*. W3C Editor's Draft 10 April 2016, 2016. Available from `http://w3c.github.io/data-shapes/shacl/`.

[12] A. Krisnadhi, P. Hitzler, and K. Janowicz. On capabilities and limitations of OWL regarding typecasting and ontology design pattern views. In V. Tamma, M. Dragoni, R. Goncalves, and A. Lawrynowicz, editors, *Ontology Engineering. 12th International Experiences and Directions Workshop on OWL, OWLED 2015, co-located with ISWC 2015, Bethlehem, PA, USA, October 9-10, 2015, Revised Selected Papers*, volume 9557 of *Lecture Notes in Computer Science*. Springer, Heidelberg, 2016. To appear.

[13] A. Krisnadhi, F. Maier, and P. Hitzler. OWL and Rules. In A. Polleres et al., editors, *Reasoning Web. Semantic Technologies for the Web of Data – 7th International Summer School 2011, Tutorial Lectures*, volume 6848 of *Lecture Notes in Computer Science*, pages 382–415. Springer, Heidelberg, 2011.

[14] M. Krötzsch. *Description Logic Rules*, volume 008 of *Studies on the Semantic Web*. IOS Press/AKA, 2010.

[15] M. Krötzsch, F. Maier, A. A. Krisnadhi, and P. Hitzler. A better uncle for

OWL: Nominal schemas for integrating rules and ontologies. In S. Sadagopan, K. Ramamritham, A. Kumar, M. Ravindra, E. Bertino, and R. Kumar, editors, *Proceedings of the 20th International World Wide Web Conference, WWW2011, Hyderabad, India, March/April 2011*, pages 645–654. ACM, New York, 2011.

[16] M. Krötzsch, S. Rudolph, and P. H. Schmitt. A closer look at the semantic relationship between datalog and description logics. *Semantic Web*, 6(1):63–79, 2015.

[17] M. Krötzsch, F. Simančík, and I. Horrocks. A description logic primer. *CoRR*, abs/1201.4089, 2012.

[18] O. Lassila and R. R. Swick. *Resource Description Framework (RDF) Model and Syntax Specification.* W3C Recommendation 10 February 2004, 2004. Available from http://www.w3.org/TR/REC-rdf-syntax/.

[19] J. W. Lloyd. *Foundations of Logic Programming.* Springer, Berlin, 1988.

[20] D. Makinson. *Bridges from Classical to Nonomonotonic Logic*, volume 5 of *Texts in Computing.* King's College Publications, 2005.

[21] D. C. Martínez, K. Janowicz, and P. Hitzler. A logical geo-ontology design pattern for quantifying over types. In I. F. Cruz, C. A. Knoblock, P. Kröger, E. Tanin, and P. Widmayer, editors, *SIGSPATIAL 2012 International Conference on Advances in Geographic Information Systems (formerly known as GIS), SIGSPATIAL'12, Redondo Beach, CA, USA, November 7-9, 2012*, pages 239–248. ACM, 2012.

[22] J. McCarthy. Circumscription – A Form of Non-Monotonic Reasoning. *Artificial Intelligence*, 13(1–2):27–39, 1980.

[23] R. Moore. Possible-worlds Semantics for Autoepistemic Logic. In *Proceedings of the 1984 Non-monotonic Reasoning Workshop.* AAAI, Menlo Park, CA, 1984.

[24] R. Reiter. A Logic for Default Reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[25] S. Rudolph, M. Krötzsch, and P. Hitzler. All Elephants are Bigger than All Mice. In F. Baader, C. Lutz, and B. Motik, editors, *Proceedings of the 21st International Workshop on Description Logics (DL2008), Dresden, Germany, May 13-16, 2008*, volume 353 of *CEUR Workshop Proceedings*, 2008.

[26] J. Tao, E. Sirin, J. Bao, and D. L. McGuinness. Integrity constraints in OWL. In M. Fox and D. Poole, editors, *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010.* AAAI Press, 2010.