An Introduction to Ontology Learning

Jens LEHMANN [a] and Johanna VÖLKER [b,1]
[a] *Informatics Institute, University of Leipzig, Germany*
[b] *Data & Web Science Research Group, University of Mannheim, Germany*

Ever since the early days of Artificial Intelligence and the development of the first knowledge-based systems in the 70s [32] people have dreamt of self-learning machines. When knowledge-based systems grew larger and the commercial interest in these technologies increased, people became aware of the knowledge acquisition bottleneck and the necessity to (partly) automatize the creation and maintenance of knowledge bases. Today, many applications which exhibit 'intelligent' behavior thanks to symbolic knowledge representation and logical inference rely on ontologies and the standards provided by the World Wide Web Committee (W3C). Supporting the construction of ontologies and populating them with instantiations of both concepts and relations, commonly referred to as *ontology learning*.

Early research in ontology learning has concentrated on the extraction of facts or schema-level knowledge from textual resources building upon earlier work in the field of computational linguistics and lexical acquisition. However, as we will show in this book, ontology learning is a very diverse and interdisciplinary field of research. Ontology learning approaches are as heterogeneous as the sources of data on the web, and as different from one another as the types of knowledge representations called "ontologies".

In the remainder of this introduction, we briefly summarize the state-of-the-art in ontology learning and elaborate on what we consider as the key challenges for current and future ontology learning research.

Ontology-based Knowledge Representation

Ontologies in computer science are usually regarded as formal representations of knowledge – often restricted to a particular domain. There is, however, no general agreement on which requirements the formal representation needs to satisfy in order to be appropriately be called an ontology. Depending on the particular point of view, ontologies can be simple dictionaries, taxonomies, thesauri, or richly axiomatized top-level formalisations. In this book, we do not limit ourselves to a particular definition of the term ontology, since use case specific requirements determine the needed expressivity of knowledge bases: the contributions range from approaches to learning lightweight structures to methods for generating complex definitions of classes.

Ontologies play a central role in data and knowledge integration. By providing a shared schema, they facilitate query answering and reasoning over disparate data sources.

[1] Johanna Völker is financed by a Margarete-von-Wrangell scholarship of the European Social Fund (ESF) and the Ministry of Science, Research and the Arts Baden-Württemberg.

Since the amount of data on the web as well as in corporate intranets, for example, is growing quickly, structured access to data becomes more important. In addition to data integration, reasoning and querying scenarios, ontologies are also a means to document the structure of a particular domain, which helps to develop a common understanding of its concepts.

However, the construction of ontologies is a highly expensive task which crucially hinges on the availability of scarce expert resources [39]. In order to build a formal ontology for a particular domain of interest, for instance, specialized domain knowledge needs to be acquired and formalized in a way that automated inference will yield the expected results. This goal can only be achieved if domain experts collaborate with skilled ontology engineers familiar with the theory and practice of knowledge representation – and once the ontology has been constructed, evolving knowledge and application requirements will demand for continuous maintenance efforts.

Ontology Learning

It is more than ten years now that Mädche and Staab [29] coined the term "ontology learning" for a newly emerging field of research aiming at nothing less than the automatic generation of ontologies. The first ontology learning workshop[2], held in 2000 and co-organized by Claire Nédellec and Peter Wiemer-Hastings, brought together people from very different research communities. Looking into the proceedings, we can distinguish works based on ripple down rules, word sense clustering, and information extraction, for example. It is remarkable that there were only few contributions from the field of concept learning at that time, although researchers have investigated the use of inductive logic programming for learning logical theories since the mid 80s. From today's perspective, ontology learning is a use case for concept learning, but the collaboration and exchange between this and other parts of the ontology learning community is still limited. Ten years after the first ontology learning workshop, we tried to compile a book which brings together the most diverse works in the area of ontology learning including contributions by the concept learning community as well as "classical" works on ontology learning from text or other semi-structured resources. It is designed to give an overview of a broad range of ontology learning approaches, logical and statistical ones, and to outline the synergies that may arise from bringing together the various methods and methodologies.

While we do not intend to draw a sharp line between different types of ontology learning, approaches can be roughly classified into the following areas:

Ontology Learning from Text mostly focuses on the automatic or semi-automatic generation of lightweight taxonomies by means of text mining and information extraction. Many of the methods used in ontology learning from text (e.g. lexico-syntactic patterns for hyponymy detection or named-entity classification) are inspired by previous work in the field of computational linguistics, essentially designed in order to facilitate the acquisition of lexical information from corpora. Some ontology learning approaches do not derive schematic structures, but focus on the data level. Such *ontology population* methods derive facts from text. A popular example is the Never-Ending Language Learning (NELL) project [10], which reads the web to add statements to its knowledge base and improves its performance over time, e.g. via user feedback.

[2] http://ol2000.aifb.uni-karlsruhe.de

Linked Data Mining refers to the process of detecting meaningful patterns in RDF graphs. One of the motivations behind this research area is that Linked Data publishers sometimes do not create an explicit schema for their dataset upfront, but focus on publishing data first. Being able to detect the structure within published RDF graphs can, on the one hand, simplify the later creation of schemata and, on the other hand, allow to detect interesting associations between elements in the RDF graph. This can be achieved via statistical schema induction [7,8,43] or statistical relational learning methods, which mine frequent patterns and correlations in large data sets. In Linked Data mining, clustering approaches can be used to group related resources and provide an enhanced structure for the underlying data.

Concept Learning in Description Logics and OWL is a direction of research that aims at learning schema axioms, such as definitions of classes, from existing ontologies and instance data. Most methods in this area are based on Inductive Logic Programming methods [33]. While many algorithms, such as DL-FOIL [16] and OCEL [25] are generic supervised machine learning approaches for description logics, there are also specific adaptations to ontology learning [22], e.g., in terms of performance and usability. Closely related to concept learning in Description Logics is *onto-relational learning*, which combines methods for learning OWL axioms with rule learning approaches [27].

Crowdsourcing ontologies is an interesting alternative to purely automatic approaches as it combines the speed of computers with the accuracy of humans. Provided that the task to be completed is simple enough, it only requires the right incentives for people to contribute. Examples of crowdsourcing in the field of ontology learning include taxonomy construction via Amazon mechanical turk, and games with a purpose for ontology population (see, e.g., [12,19]).

Other approaches include, e.g., transfer learning and ontology re-use, which try to adapt existing ontologies to new domains by partially re-using existing schematic structures. Furthermore, apart from the above mentioned combination of rules and ontologies, direct representations of uncertainty, e.g. via Markov Logic Networks, are also investigated.

Major and minor distinctions between these approaches make it difficult to come up with a formal definition [41], or a breakdown into concrete subtasks (see, e.g., the *ontology learning layer cake* [35]) which is neither too general nor limited to one particular type of approach [15,29,45]. However, it is this variety that makes the ontology learning community so rich and inspiring. A lot of progress has been made in each of the above-mentioned fields, and we can observe a trend towards hybrid and integrated approaches. For this book, we assembled contributions by researchers addressing some of the key challenges in ontology learning:

Heterogeneity. Data on the web differs largely, e.g., with respect to formats, languages, domains and quality.[3] Approaches to learning from heterogeneous sources of evidence [9] can effectively leverage this huge variety by increasing the accuracy as well as the coverage of learned ontologies. However, neither the integration of methods nor the homogenization of data has attracted high attention within the ontology learning commu-

[3] See [28] for an early attempt to categorize the different types of data that can serve as input to ontology learning approaches.

nity so far and remains to be a challenge for the application of ontology learning methods in practice.

Uncertainty. Low-quality or unstructured data, which is hard to interpret by computational means, as well as inherently imperfect methods for learning ontologies can lead to results that are less likely to be correct. Thus, many ontology learning methods are designed in a way that they associate each individual outcome with a certainty value reflecting the methods' *confidence* with regard to the correctness of particular results. Uncertainty values and other types of provenance information such as timestamps or authorship annotations are especially important when it comes to manual or automatic debugging of learned ontologies [14].

Reasoning. Often, ontologies are learned or manually created for applications which are based on logical inference. In case these applications require the ontology to be logically consistent, ontology learning approaches should be capable of generating consistent (and coherent) ontologies. Therefore, not only methods for concept learning in description logics, but also other ontology learning approaches rely on logical reasoning. Experiments have shown that a tight coupling between ontology debugging, i.e. inconsistency diagnosis and repair, and ontology learning may be beneficial [31,23].

Scalability. Extracting knowledge from the growing amounts of data on the web – unstructured, textual data on the one hand and structured data such as databases, linked data[4] or ontologies on the other hand – requires scalable and efficient approaches. Especially when it comes to learning from distributed, loosely interconnected data and the integration of knowledge from multiple sources, ontology learning methods face big challenges. In order to address these challenges, various strategies are currently being developed, such as distributed computation for horizontally scaling ontology learning, incremental learning approaches for re-using existing knowledge, or sampling [17] and modularization to improve the efficiency of ontology learning algorithms.

Quality. We can measure the quality of an automatically generated ontology as we can measure the quality of any ontology, be it learned or manually engineered. However, *ontology evaluation* is not an easy task (see [44] for a comprehensive overview of the state-of-the-art in ontology evaluation). Formal correctness, completeness and consistency are only a few of many possible criteria for judging the quality of an ontology, and it is the application context of an ontology which ultimately determines the choice of evaluation criteria as well as the required quality standard. Ideally, each step of an ontology learning process, including the choice of input data as well as preprocessing and relation extraction, for example, should thus be optimized with regard to the particular domain or application context that the learned ontology will be used for. Ontology learning methodologies [40] and the adoption of ontology design patterns [4] can help to further improve the results.

Interactivity. In practice, the quality of a learned ontology often depends on the degree of automation. The lesser the extent to which humans are involved in a semi-automatic ontology generation process, the lower the quality we can expect. An ontology insufficient for the intended application in terms of quality, e.g., after having been generated in a fully automatic way, will eventually require a significant amount of post-processing.

[4]See `http://stats.lod2.eu` and `http://lod-cloud.net`.

While the revision of learned ontologies is not generally considered part of the actual ontology learning process, methods for automatic ontology generation can support this sort of post-processing by providing detailed provenance information. Provenance information acquired in the course of ontology learning typically comprises, for example, confidence and relevance values for individual axioms as well as time stamps and key facts about the employed learning procedure. Nevertheless, the amount of post-processing can be a significant burden for knowledge engineers, and innovative methods are required in order to overcome the so-called *knowledge acquisition bottleneck*. Generally, it is advisable to integrate methods for ontology learning and revision into popular ontology engineering frameworks, in order to reduce the overhead for human participation in the overall process [22,37]. Crowdsourcing and games with a purpose can help to lower the costs of revising learned ontologies by involving non-experts, but translating their interactions into ontology modeling decisions is a non-trivial problem. Systematic expert interrogation, known as relational exploration, has been found to be an efficient way of asking people the right questions, while at the same time reducing the overall number of decisions to be made [2]. Finally, experiments have shown that ontology design patterns, which capture the knowledge and experience of human ontology engineers, can beneficially be integrated into the ontology learning process [5].

About this Book

In the following, we describe the structure of the book and give an outline of the content of individual chapters.

Foundations This part of the book covers the most basic concepts which are important to understand state-of-the-art ontology learning approaches. First, Krötzsch et al. [20] give a primer on description logics, the family of knowledge representation formalisms which underly most of the learned ontologies. The second chapter contributed by Jentzsch, Vrandečić and Usbeck [18] introduces RDF and the Linked Data principles as well as fundamental semantic web technologies such as SPARQL, for example. It is followed by an overview of typical machine learning approaches which are applied in ontology learning (see Ławrynowicz and Tresp [21]). Finally, as the vast majority of ontology learning algorithms still relies on textual input, Maynard and Bontcheva [30] cover the most relevant natural language processing techniques.

Logical Learning The second part of this book focuses on approaches which have been developed to derive ontologies from structured knowledge. Logical learning methods are introduced by Lehmann et al. [24], who cover the foundations of learning description logic concepts, as well as by Francesca Lisi [26] who contributed a chapter on learning onto-relational rules. The presented methods are based on the assumption that schema axioms, such as concept definitions, can be learned from existing instance level knowledge.

Lexical Learning Given the vast amounts of textual contents on the web it is not astonishing that previous research in ontology learning mainly concentrated on the acquisition of ontologies from unstructured data. In this part of the book, we therefore present contributions to ontology learning from natural language text based on information extraction and text mining techniques. The first chapter (cf. Coppola et al. [13]) introduces a methodology for generating domain-specific ontologies by specializing and instantiating frames from the FrameNet lexicon. In the second chapter, Fabian Suchanek [42] gives an overview of well-known methods for acquiring facts and taxonomies from Wikipedia articles, while putting an emphasis on the synergies between ontological reasoning and information extraction. Finally, the chapter contributed by Nováček and Handschuh [34] outlines a layered ontology learning framework, which facilitates the integration of facts extracted from textual documents with existing schema-level knowledge.

Learning from Web Data The increasing amounts of data available through the web and the growing need for automatic approaches to the access and usage of information on the web, pose particular challenges to ontology learning methods. This part of the book is therefore dedicated to methods which take into account the specific characteristics of web data, such as heterogeneity, volume, decentralization and a large variety of different formats. A part of the world wide web, which has seen a tremendous rise in importance over the past decade, are social websites and crowdsourced tagging applications. The first chapter in this part of the book gives a systematic overview over the trends and future developments in the area of knowledge extraction from tagging systems or folksonomies (see Benz and Hotho [3]). Spatial applications and annotations of points of interests are investigated by Alves and Pereira [1] in the second chapter, which deals with approaches to semantically enrich the descriptions of spatial objects. Finally, in the last chapter in this part, Cerbah and Lammari [11] address the problem of web application backends. Motivated by the fact that the vast majority of web applications are driven by relational databases, they present methods for deriving ontologies from schemata of relational databases.

Dynamics and User Interaction In many cases, purely automatic learning approaches will fail to generate ontologies which are good enough for a particular, e.g. reasoning-based, application. For this reason, we would like to emphasize the role users and knowledge engineers can play in an ontology learning process. The first chapter of this part written by Simperl et al. [38] explains how games with a purpose can help to leverage human resources for learning ontologies. A completely different approach to "putting the human in the loop" is presented by Rudolph and Sertkaya [36], who introduce formal concept analysis as a means to interactive ontology learning and as an effective way to minimize the required human effort. Last, but not least, Blomqvist et al. [6] elaborate on key challenges in ontology learning and how those can be addressed by the use of ontology design patterns, which encode best-practices in ontology engineering.

References

[1] Ana Oliveira Alves and Francisco Camara Pereira. Semantic enrichment of places: From public places descriptions to linked data. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[2] Franz Baader, Bernhard Ganter, Barış Sertkaya, and Ulrike Sattler. Completing Description Logic knowledge bases using Formal Concept Analysis. In Manuela M. Veloso, editor, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 230–235, 2007.

[3] Dominik Benz and Andreas Hotho. Capturing emergent semantics from social tagging systems. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[4] Eva Blomqvist. Ontocase – a pattern-based ontology construction approach. In Robert Meersman and Zahir Tari, editors, *On the Move to Meaningful Internet Systems 2007: CoopIS, DOA, ODBASE, GADA, and IS*, volume 4803 of *Lecture Notes in Computer Science*, pages 971–988. Springer Berlin Heidelberg, 2007.

[5] Eva Blomqvist. *Semi-automatic Ontology Construction based on Patterns*. PhD thesis, Linköping University, Department of Computer and Information Science at the Institute of Technology, 2009.

[6] Eva Blomqvist, Aldo Gangemi, and Francesco Draicchio. Ontology design patterns in ontology learning. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[7] Lorenz Bühmann and Jens Lehmann. Universal OWL axiom enrichment for large knowledge bases. In *Proceedings of EKAW 2012*, pages 57–71. Springer, 2012.

[8] Lorenz Buhmann and Jens Lehmann. Pattern based knowledge base enrichment. In *12th International Semantic Web Conference, 21-25 October 2013, Sydney, Australia*, 2013.

[9] Paul Buitelaar, Philipp Cimiano, Anette Frank, Matthias Hartung, and Stefania Racioppa. Ontology-based information extraction and integration from heterogeneous data sources. *Journal on Human-Computer Studies*, 66(11):759–788, 2008.

[10] A. Carlson, J. Betteridge, B. Kisiel, B. Settles, E.R. Hruschka Jr, and T.M. Mitchell. Toward an architecture for never-ending language learning. In *Proceedings of the 24th Conference on Artificial Intelligence (AAAI)*, volume 2, pages 1306–1313, 2010.

[11] Farid Cerbah and Nadira Lammari. Ontology learning from databases: Some efficient methods to discover semantic patterns in data. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[12] Lydia B. Chilton, Greg Little, Darren Edge, Daniel S. Weld, and James A. Landay. Cascade: Crowdsourcing taxonomy creation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 1999–2008, New York, NY, USA, 2013. ACM.

[13] Bonaventura Coppola, Aldo Gangemi, Alfio Gliozzo, Davide Picca, and Valentina Presutti. Learning domain ontologies by corpus-driven framenet specialization. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[14] Paulo Cesar G. da Costa, Kathryn B. Laskey, Kenneth J. Laskey, and Michael Pool, editors. *International Semantic Web Conference, ISWC 2005, Galway, Ireland, Workshop 3: Uncertainty Reasoning for the Semantic Web, 7 November 2005*, 2005.

[15] X.Y. Du, M. Li, and S. Wang. A survey on ontology learning research. *Journal of Software*, 17(9):1837–1847, 2006.

[16] Nicola Fanizzi, Claudia d'Amato, and Floriana Esposito. DL-FOIL: Concept learning in description logics. In F. Zelezný and N. Lavrac, editors, *Proceedings of the 18th International Conference on Inductive Logic Programming (ILP)*, volume 5194 of *LNAI*, pages 107–121. Springer, 2008.

[17] Sebastian Hellmann, Jens Lehmann, and Sören Auer. Learning of OWL class descriptions on very large knowledge bases. *International Journal on Semantic Web and Information Systems*, 5(2):25–48, 2009.

[18] Anja Jentzsch, Ricardo Usbeck, and Denny Vrandečić. An incomplete and simplifying introduction to linked data. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[19] Dimitris Karampinas and Peter Triantafillou. Crowdsourcing taxonomies. In *Proceedings of the 9th International Conference on The Semantic Web: Research and Applications*, ESWC'12, pages 545–559, Berlin, Heidelberg, 2012. Springer-Verlag.

[20] Markus Krötzsch, Frantisek Simančík, and Ian Horrocks. A description logic primer. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[21] Agnieszka Ławrynowicz and Volker Tresp. Introducing machine learning. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[22] Jens Lehmann, Sören Auer, Lorenz Bühmann, and Sebastian Tramp. Class expression learning for ontology engineering. *Journal of Web Semantics*, 9:71 – 81, 2011.

[23] Jens Lehmann and Lorenz Bühmann. ORE – a tool for repairing and enriching knowledge bases. In *Proceedings of the 9th International Semantic Web Conference (ISWC)*, Lecture Notes in Computer Science, Berlin / Heidelberg, 2010. Springer.

[24] Jens Lehmann, Nicola Fanizzi, and Claudia d'Amato. Concept learning. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[25] Jens Lehmann and Pascal Hitzler. Concept learning in description logics using refinement operators.

Machine Learning journal, 78(1-2):203–250, 2010.

[26] Francesca Lisi. Learning onto-relational rules with inductive logic programming. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[27] Francesca A. Lisi and Floriana Esposito. Nonmonotonic onto-relational learning. In Luc De Raedt, editor, *ILP*, volume 5989 of *Lecture Notes in Computer Science*, pages 88–95. Springer, 2009.

[28] Alexander Mädche. *Ontology Learning for the Semantic Web*. PhD thesis, Universität Karlsruhe (TH), Germany, 2001.

[29] Alexander Mädche and Steffen Staab. Ontology learning for the semantic web. *IEEE Intelligent Systems*, 16(2):72–79, 2001.

[30] Diana Maynard and Kalina Bontcheva. Natural language processing. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[31] Christian Meilicke, Johanna Völker, and Heiner Stuckenschmidt. Learning disjointness for debugging mappings between lightweight ontologies. In Aldo Gangemi and Jerome Euzenat, editors, *Knowledge Engineering: Practice and Patterns*, volume 5268 of *Lecture Notes in Computer Science*, pages 93–108. Springer Berlin Heidelberg, 2008.

[32] M. Minsky. A framework for representing knowledge. In P. Winston, editor, *The Psychology of Computer Vision*, pages 211–277. McGraw-Hill, New York, 1975.

[33] Shan-Hwei Nienhuys-Cheng and Ronald de Wolf, editors. *Foundations of Inductive Logic Programming*, volume 1228 of *Lecture Notes in Computer Science*. Springer, 1997.

[34] Vit Novacek and Siegfried Handschuh. Empirically grounded emergent knowledge. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[35] Bernardo Magnini Paul Buitelaar, Philipp Cimiano. *Ontology Learning from Text: Methods, Evaluation and Applications*. IOS Press, 2008.

[36] Sebastian Rudolph and Baris Sertkaya. Formal concept analysis methods for interactive ontology learning. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[37] Barış Sertkaya. OntoComP: A protégé plugin for completing OWL ontologies. In *The Semantic Web: Research and Applications*, volume 5554 of *LNCS*, pages 898–902. Springer Berlin Heidelberg, 2009.

[38] Elana Simperl, Stephan Wölger, Stefan Thaler, and Katharina Siorpaes. Learning ontologies via games with a purpose. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[39] Elena Simperl, Tobias Buerger, Simon Hangl, Stephan Woelger, and Igor Popov. Ontocom: A reliable cost estimation method for ontology development projects. *Web Semantics: Science, Services and Agents on the World Wide Web*, 16(0):1 – 16, 2012.

[40] Elena Simperl, Christoph Tempich, and Denny Vrandečić. A methodology for ontology learning. In *Proceedings of the Workshop on Ontology Learning and Population: Bridging the Gap between Text and Knowledge*, pages 225–249. IOS Press, 2008.

[41] Michael Sintek, Paul Buitelaar, and Daniel Olejnik. A formalization of ontology learning from text. In *Proceedings of the Workshop on Evaluation of Ontology-based Tools (EON) at the International Semantic Web Conference (ISWC)*, 2004.

[42] Fabian Suchanek. Information extraction for ontology learning. In Jens Lehmann and Johanna Völker, editors, *Perspectives on Ontology Learning*, Studies on the Semantic Web. AKA Heidelberg / IOS Press, 2014.

[43] Johanna Völker and Mathias Niepert. Statistical schema induction. In Grigoris Antoniou, Marko Grobelnik, Elena Paslaru Bontas Simperl, Bijan Parsia, Dimitris Plexousakis, Pieter De Leenheer, and Jeff Z. Pan, editors, *The Semantic Web: Research and Applications - 8th Extended Semantic Web Conference (ESWC), Heraklion, Crete, Greece, May 29-June 2, 2011, Proceedings*, volume 6643 of *Lecture Notes in Computer Science*, pages 124–138. Springer, 2011.

[44] Denny Vrandečić. Ontology evaluation. In Steffen Staab and Rudi Studer, editors, *Handbook on Ontologies*, International Handbooks on Information Systems, pages 293–313. Springer, 2009.

[45] W. Wong, W. Liu, and M. Bennamoun. Ontology learning from text: A look back and into the future. *ACM Computing Surveys (CSUR)*, 44(4):20, 2012.