Comparison of Inductive Inference Mechanisms and their Suitability for an Information Model for the Visualization of Uncertainty

Moritz Weber^{1,a*}, Georg Staudter^{1,b}, Reiner Anderl^{1,c}

¹Otto-Berndt-Straße 2, 64287 Darmstadt, Germany

^am.weber@dik.tu-darmstadt.de, ^bstaudter@dik.tu-darmstadt.de,

^canderl@dik.tu-darmstadt.de

Keywords: Uncertainty, Ontologies, OWL, Inductive Reasoning

Abstract. Ontologies represent inter-related semantic information. The automated integration of new knowledge helps to detect and reduce data-induced conflicts and model-based uncertainty in ontologies. However, automatic extension of an existing ontology from heterogeneous distributed sources can often lead to incomplete and contradictory entities. In order to resolve these conflicts and to complete entities, inductive inference mechanisms should be applied in addition to the deductive mechanisms already in use. This paper first describes various inductive inference mechanisms and compares these with each other according to pre-defined requirements and other criteria. Finally, the mechanisms' suitability for an information model for the exchange and visualization of uncertainty in load-carrying systems and possible combinations of the individual mechanisms are discussed, also with respect to the necessity of further modifications of these mechanisms.

Introduction

To detect and reduce uncertainty in load-carrying systems and thereby assist the designing engineer by providing valuable information, it is advantageous to include as much data from as many sources as possible. In the context of this paper, this means to extend an existing ontology with information from heterogeneous and distributed sources. The integration of this information inevitably leads to incorrect and incomplete data. Reasoning can be used to infer new information based on existing knowledge from an ontology or to determine inconsistencies because the existing knowledge is contradicting. Reasoning in general means the drawing of conclusions on the basis of principles and evidence [1,2]. There are two types of reasoning: Inductive and deductive reasoning. In deductive reasoning new knowledge is acquired by using more general knowledge ("top-down" e.g. the axioms "birds are animals" and "a dove is a bird" implies that a dove is an animal) whereas in inductive reasoning one or more assets are used to find a more general rule ("bottom-up" e.g. the finding that all known birds have feathers leads to the assumption that all birds have feathers so that a dove presumably also has feathers).

This leads to a significant difference in the certainty of the conclusions: As the general rules in deduction have to be universal, the conclusions found are either valid or invalid. By contrast, the rules found by induction are inherently uncertain. There is no way to certainly acquire a rule only based on given examples.

In ontologies every entity needs to be assigned to at least one predefined class, the so-called concepts. For ontologies it is therefore possible to infer concept membership with certainty with the help of deductive reasoning and additionally to check the integrity of the entire ontology. For deductive reasoning several so-called semantic reasoners have been developed and are constantly being improved. A semantic reasoner is a software that is able to infer logical consequences from a knowledge base. This means that current reasoners can infer concept memberships as well as relationships between entities, which are logically deductible. Some widely-known examples especially for ontologies are Pellet [3] and HermiT [4]. However, in order to obtain the most information possible from an ontology, information should additionally be obtained by inductive reasoning. There is currently no reasoner available for inductive reasoning. With regard to

ontologies, inductive reasoning means in particular the classification of a new asset based on existing assets, where a classification solely based on deduction is not possible. In this paper, different methods for inductive reasoning in ontologies are discussed and compared.

The ontological information model developed within the Collaborative Research Centre 805 (SFB 805) shows the relationships between uncertain geometries and components. In the current third funding period, an automatic extension of the ontology is to be made possible from distributed and heterogeneous sources. For this, two new approaches are to be pursued. Ontology matching [5] on the one hand to integrate other ontologies or parts of ontologies and inductive reasoning to acquire new knowledge on the other hand. Applying inductive reasoning in addition to deductive reasoning also helps to obtain the largest possible data basis, leading to more elaborate results. Moreover, it is possible to detect inconsistencies requiring manual resolution. By this means, a larger amount of knowledge can be used to reduce uncertainty, which is especially beneficial in the context of load-carrying systems.

Structure of the Paper

This paper compares different inductive inference mechanisms mentioned in literature with regard to their applicability in an ontology-based information model for the exchange and visualization of load-carrying systems which is being developed in the SFB 805 at Technische Universität Darmstadt. In the section background some fundamentals of ontologies necessary to understand the remainder of the paper are explained. The basic structure and use of the information model are then described. This results in special requirements. Based on the previous sections, various approaches mentioned in literature are summarised and their advantages and disadvantages are discussed afterwards. The degree to which the requirements stated before are met is also discussed. The consequences from the discussion are then summarised in the results section and a possible combination of the individual methods is proposed. Finally, an outlook on how to integrate the algorithms into the information model and a brief conclusion is given.

Background

Description Logics (DL) are a group of Languages for knowledge representation e.g. in ontologies. For simplicity, we will focus on ALC [6], which is one of the most basic description logics and is supported by most ontology languages. In ALC, descriptions are defined by three disjoint sets of names N_C , N_R , and N_I . They denote concepts (e.g. classes like *building*), roles (e.g. relations like *standsIn*) and individuals (e.g. instances like *EiffelTower*). The top concept $\top$ (also referred to as *thing*) is the generalization of all classes whereas the bottom concept $\perp$ (also referred to as *nothing*) is the empty concept.

Concepts can be defined as an atomic concept in N_C or as different expressions. Complement, Union and Intersection are represented by the operators $\neg$, $\sqcup$ and $\sqcap$. The expression $\exists R.C$ denotes the set individuals, who are related through a certain Role R to at least one Instance of the Concept C . The expression $\forall R.C$ denotes the set of individuals, who are related through R only with instances of C . For example the Expression $\exists \text{standsIn}. \text{City}$ describes the set of individuals who stand in a city whereas $\forall \text{hasResident}. \text{MalePerson}$ is the set of all individuals, which have only male Persons as their inhabitants.

A knowledge base K comprises of the TBox T and the ABox A , so that $K = \langle T, A \rangle$. The TBox is a set of definitions in the form of $C \sqsubseteq D$, where C and D are concept expressions. This axiom expresses that every instance of D is also an instance of C . For example the expression $\text{CityBuilding} \sqsubseteq \text{Building} \sqcap \exists \text{standsIn}. \text{City}$ means that every entity in K that is an instance of *building* and stands in an instance of *City* is also a *CityBuilding*. The ABox consists of assertions in the form $A(a)$ and $R(a,b)$ where $A \in N_C$, $R \in N_R$ and $a, b \in N_I$. Thereby is $A(a)$ the assertion of an individual to a concept and $R(a,b)$ defines the relationship between two concepts.

Ontologies are a special kind of knowledge bases [7], which are meant to be an “explicit specification of a conceptualization” (Gruber, 1993) [8] or a “formal specification of a shared conceptualization” (Borst, 1997) [9]. Ontologies are used to describe the semantics of a domain of interest. Therefore, they use the extensions of ALC to formalize the available knowledge [10]. One of the most important Ontology Languages is the Web Ontology Language (OWL) [11]. It was designed by the World Wide Web Consortium (W3C) and uses the Resource Description Framework (RDF) and XML as their foundation.

In RDF, expressions consist of triples of subject, predicate and object, where the subject is a class (in the TBox) or an Instance of a class (in the ABox), the predicate corresponds to a role in DLs and the object is again a class or instance but can also be a datatype (e. g. such as integer or string). OWL extends the capabilities of ALC and RDF for example by expressions to describe that two classes (or concepts in DL) are disjoint or that two relations are the inverse of each other (e. g. *isPartOf* is the inverse to *hasPart*). There are three styles of OWL which are called OWL Full, OWL DL and OWL Lite. OWL Full has full upward compatibility with RDF but is not decidable, which restricts the use of a semantic reasoner. OWL DL is the decidable subset OWL Full. OWL Lite is then again subset of OWL DL and is designed to be a much simpler and lighter version of OWL DL.

Opposing to classical databases, which use the closed world assumption (CWA), ontologies use the open world assumption (OWA). The OWA denotes that a statement cannot be assumed to be true if its negation cannot be proven (e. g. you cannot know if somebody has no children only because he has no children defined in the ontology). That is a big advantage in the field of uncertainty regarding completeness of data, because no wrong assumptions will be made due to missing data.

Requirements by the Ontology-Based Information Model for an Inductive Reasoning Mechanism

In the first funding period of the SFB 805, an ontology-based information model for the exchange and visualization of uncertainty in load-carrying systems has been developed. This model allows to analyse the interactions of uncertainties of individual components and the entire product and for displaying their effects to the designer [12]. In the second funding period of the SFB 805, it was extended to include the possibility of representing time-dependent uncertainty data [13]. In the current third period an automatization of the integration of additional data is to be done.

Figure 1. Visualization of an extract of the ontology developed in SFB 805

The TBox of the ontology developed in the first two phase of SFB 805 is based on ISO standard 10303-108 and the ABox consists of ontological representations of explicit CAD models generated by OntoStep [14]. These models are described in the ontology by their geometric and topologic data

and also data on uncertainty. Figure 1 shows a part of an example of the ontology developed. Here the concepts are the entities with a circle and the instances with a diamond besides the name. The arrows show different relations and assertions.

This paper examines algorithms for inducing new knowledge on knowledge bases and their applicability for the information model developed in SFB 805. This results in special requirements to be met by the algorithms, which will be explained in the following.

Since the information model used is intended to support designers, they should be provided with a measure of the reliability of the statements made by induction in order to be in a position to deal with uncertainties and, if necessary, to take appropriate countermeasures. Ideally, the reasons of the conclusions should be clearly identifiable in order to take well-informed decisions. This means that the designer can see the information (i.e. the entities and axioms) used for the conclusion.

As the information model is developed in OWL DL, the mechanism should ideally have been designed for this language or at least be transferable. However, it should also be noted that this can lead to a change in the extent to which the results described in the literature are still be valid.

If the mechanisms have further requirements concerning the underlying ontology, these will be discussed in the corresponding section.

Mechanisms Proposed in Literature

Fanizzi et al. [15] present a classification method using parametric language-independent kernel functions. Kernel functions are a class of efficient algorithms for pattern analysis [16]. They employ so-called kernels to operate implicitly in a higher-dimensional feature space. Hereby a non-linear pattern analysis can be transformed to a linear one. One of the known algorithms utilizing kernel functions is for example Support Vector Machines (SVM), which are also used by the proposed method for classification. The authors propose the algorithm especially for supervised population of the ABox i.e. that the classification should be confirmed manually.

For the evaluation purposes, five OWL-ontologies that are written in different DL-languages are used. The ontologies contain 19 to 112 concepts and 72 to 676 individual entities. The algorithm is validated against a deductive reasoner, so that the match rate is the percentage of classifications done by both the deductive reasoner and the proposed algorithm. The induction rate is the ratio between classifications done by the method not logically deductible and all classifications done by the method. The results show match rates between 82 % and 98 %. The authors state that the match rate increases with the number of entities in the considered ontology. The induction rate was between 0 and 2%.

The used kernel method makes it impossible to understand the reasons for a single classification, but it is generally possible to provide a measure for the correctness of a classification.

D'Amato et al. [17,18] propose two similar methods of instance-based learning for inductive reasoning on ABoxes. These can be used for classification, but also for predicting new assertions for the ontology. They both use a variation of the k-Nearest-Neighbours (k-NN) algorithm that uses the k nearest neighbours to vote for the class to be assigned. To determine the nearest neighbours, two different composite dissimilarity measures for ALC are used. The first method uses a measure based on overlap of the properties of the entities, whereas the other method is based on the information given by the entities. These are described in more detail in [19]. They both use the structure as well as the semantics of the entities.

The algorithms were validated on four different ontologies, all of which were realized in OWL DL. However, the ontologies differ considerably in the underlying description logic and the size and heterogeneity of the entities. The experiments show that the results improve significantly with increasing size and heterogeneity of the data base. This applies to both measures, as well as to both use cases, the classification and the prediction of new assertions. The smallest average match are 60.8 % for the method employing the measure based on information and 65,4 % for the method employing the measure based on overlap respectively. The highest match rate is 80.7 % for the largest and most heterogeneous ontology. The results for the larger two of the four ontologies do not differ between the two dissimilarity measures.

A variation of this method [20] employs an entropic distance measure. In contrast to the dissimilarity measures mentioned above, this measure does not use the structural criteria, so that it can deal with all standard ontology languages. The modified algorithm was validated using six OWL ontologies. The results for the match rate are between 93.3 % and 99.5 % and the induction rate is between 0 and 4.2 %.

By using the k-NN algorithm, it is possible to indicate the reasons for the classification, namely the nearest neighbours, and also to provide a measure for the correctness by indicating the proximity to the nearest neighbours.

Bourauoui et al. [21] use so-called conceptual spaces (CS) for ABox and TBox Reasoning on OWL ontologies. CS use a special form of Bayesian inference over interpretable feature representations. This allows for both seeing the base for the inference and obtaining a measure of correctness for each conclusion. CS are geometric representations using so-called quality dimensions [22]. These can be characteristics such as weight, colour, age or the three spatial dimensions. Quality dimensions make it possible to represent an entity as a vector of its properties. To transfer the entities of an ontology from the original feature space (i.e. the space used in the given ontology) into conceptual space more general taxonomies and/or ontologies are needed. Bourauoui et al. apply a method that uses textual descriptions and structured information from WikiData and WordNet as upper ontologies [23]. Upper ontologies (also called top-level ontologies) are ontologies, which try to formalize the knowledge about general terms that are valid across all domains [24].

The validation is performed using SUMO [25,26], a large freely available ontology. SUMO uses an ontology language, especially developed for it. However, a translation of SUMO to OWL is provided the developers that involves a loss of data. This translation is used for validation.

The proposed method is compared to a simple similarity based reasoning method as a baseline i. e. using the most similar entity for classification. The results show a better performance of the proposed method on all test sets, especially in terms of precision for concepts with many instances. This applies to both TBox and ABox reasoning.

Another method using conceptual spaces is proposed by Derrac and Schockaert [27]. In contrast to the method used by Bourauoui et al., the reasoning is not only based on similarities but also with a fortiori (Latin: from the stronger) reasoning and so-called Betweenness. Betweenness here means that the algorithm tries to find pairs of entities belonging to the same concept such that the entity to be classified lies approximately between this pair. If the algorithm finds more than one pair, the majority votes for the class to be asserted. For the conversion to CS the authors utilize the upper ontology OpenCyc and the taxonomies from GeoNames and Foursquare.

The algorithm is validated against SVM and k-NN algorithms on three different knowledge bases generated by the authors themselves. The results show that the classifier is able to outperform the SVM and k-NN algorithms. This is especially valid when using smaller knowledge bases, where similarity-based reasoning has systematic disadvantages.

The method used makes it easy to understand the reasons for the classification of entities given the pairs of entities used. The measures for correctness of classification on the other hand are hard to obtain and only locally comparable. Even though the method is not designed for or validated on OWL ontologies it is possible to transfer it to these as conceptual spaces are used.

Adaptive Knowledge Propagation with Dissimilarities (AKP^D) [28] uses homophilic and heterophilic relations to predict missing properties and class assignments. Homophilic relations are those types of relations that connect mainly similar instances and heterophilic relations are those that connect opposing entities (e.g. people linked by *friendWith* often live in the same country, whereas people linked by *marriedWith* mostly have different genders). AKP^D is divided into two steps: In the learning step the homophilic and heterophilic relations are identified, which are then used in the inference step for the classification and prediction of properties of entities. In the learning step, an optimization is necessary to find the two types of relations using a defined cost function.

The validation of the method is carried out using four ontologies written in different description logics. The method is compared to three other methods proposed for learning from semantic web knowledge bases. The results show a larger area under curve for the used precision/recall graphs.

For all predictions the probabilities for the correctness of the results are calculated and via the homophilic and heterophilic links it is also possible to recognize the origin of these statements.

Discussion

Table 1 shows a summary of the most important properties of the previously presented methods for inductive reasoning. It is obvious that the different variations of k-NN are the only ones that fully meet the requirements described above. They do not require any further data, are easily applicable to different ontologies and provide both a measure of correctness of the conclusions and the ability to comprehend the reasons for a classification. However, the results of Minervini et al. [28] and Derrac and Schockaert [27], in particular show that k-NN algorithms are outperformed in match rate or precision by other methods. This is especially true for small ontologies, as Derrac and Schockaert show.

A comparison of the quality of the individual algorithms by means of the correctness of the classification is impossible, since the authors use both different quality measures (e.g. match rate, precision, F1 score and area-under-curve). A comparison would also require that all validations were performed on the same ontologies.

The methods that use CS are unsuitable for the ontology of the SFB 805, since an ontology is needed that defines the terms used in the proposed ontology. However, there is no such ontology in the field of uncertainty yet and general upper ontologies cannot be applied because they are not specific enough in terms of uncertainty. Upper ontologies are not applicable because they do not define the terms used in this context detailed enough. The Betweenness suggested by Derrac and Schockaert could, however, also be transferred to the original feature space, making an upper ontology unnecessary. To which extent the results of this method presented above are transferable needs further investigation.

Methods that use kernel functions seem unsuitable for the intended use, since their application hides the reasons for conclusions through the transformation of the feature space, even if a measure of the correctness of the statements is possible. On the other hand, this class of algorithms is very efficient and gives promising results in terms of classification quality, so they should receive further attention.

Table 1. Overview of the discussed methods

Source	Used mechanism	Additional data needed	Measurement of correctness	Reasonable conclusions	Validated in OWL DL
[15]	Kernel (SVM)	no	yes	no	yes
[17]	k-NN (dissimilarity)	no	yes	yes	yes
[20]	k-NN (entropic dist.)	no	yes	yes	yes
[21]	NN (CS)	yes (CS)	yes	yes	yes
[27]	NN and Betweenness (CS)	yes (CS)	no	yes	no
[28]	Relations	no	no	yes	yes

Results

The discussion shows, that there are methods, which can be adapted for the information model developed in SFB805 without fundamental changes. But the results also show, that there is no superior method in all situations and conditions. Therefore a hybrid algorithm that combines several of the aforementioned algorithms is taken into consideration. This Combination aims to unite advantages and eliminate limitations of the individual methods.

Since the classification quality of k-NN based algorithms depends significantly on the number of entities, the utilization of Betweenness proposed by Derrac and Schockaert could be beneficial to make trustworthy classifications even if the number of entities is relatively low. To this end, a method to detect when there is too limited data for inductive reasoning with k-NN algorithms would have to be developed.

A combination of the methods described above with the method proposed by Fanizzi et al. [13] could exploit its high quality in classification by including it in the classifier. In this case, however, the reasons for the inclusion would still be represented based on the classification made by k-NN or Betweenness, even if a part of the causes of conclusions is not comprehensible this way.

Outlook

Further investigations should be carried out taking a closer look at the algorithms the different methods are based on. Possible implementations for the information model developed by SFB 805 should be planned and performed. In a next step these implementations should be tested in different scenarios. In this manner a comparison of the quality of the conclusion will be possible. The results of the tests can then be used to examine possible combinations of the algorithm in a new concept and compared this to the algorithms presented in this paper. Finally a concept to combine the inductive reasoning algorithm with ontology matching and to integrate the results in the existing visualization needs to be elaborated and implemented. The result is a tool to provide design engineers the information on uncertainty from different results so they can better control uncertainty in designing their systems.

Conclusion

This paper explained the basics of inductive reasoning, description logics and ontologies. The structure of the ontology developed in SFB 805 is presented and the related requirements are given. Since the method to be developed, should be used by designers of load-carrying systems they should be able to assess the quality of the classifications. Afterwards, several methods for classifying entities inductively for OWL ontologies mentioned in literature were discussed. The methods use different underlying principles and have different advantages and disadvantages. The differences between the algorithms and their suitability for the ontology developed for SFB 805 were discussed. It can be shown that there is no method, which is likely to outperform all other methods in every setting. Methods based on k-NN algorithms fulfil all requirements, but show weaknesses in sparsely populated ontologies. Therefore, a concept to combine the advantages of different methods is briefly proposed, which combines k-NN with other algorithms improve the classifications. In the future, this combinational concept is to be elaborated, implemented and tested on the information model developed in SFB 805 and a concept to visualize the results needs to be developed.

Acknowledgements

The authors like to thank the German Research Foundation DFG for funding this research within the Collaborative Research Centre (SFB) 805 “Control of Uncertainties in Load-Carrying Structures in Mechanical Engineering”

References

- [1] Wason P. C., and Johnson-Laird P. N., 1975. *Psychology of reasoning: Structure and content*, 2nd ed., Harvard University Press, Cambridge, Mass. [u.a.].
- [2] Sternberg R. J., 2009. *Cognitive Psychology*, Wadsworth, Belmont, CA.
- [3] Sirin E., Parsia B., Grau B. C., Kalyanpur A., and Katz Y., 2007, "Pellet: A practical OWL-DL reasoner," *Web Semantics: Science, Services and Agents on the World Wide Web*, **5**(2), pp. 51–53.
- [4] Glimm B., Horrocks I., Motik B., Stoilos G., and Wang Z., 2014, "HermiT: An OWL 2 Reasoner," *J Autom Reasoning*, **53**(3), pp. 245–269.
- [5] Euzenat J., and Shvaiko P., 2007. *Ontology matching*, Springer, Berlin u.a.
- [6] Baader F., Ed., 2003. *The description logic handbook: Theory, implementation, and applications*, 1st ed., Cambridge Univ. Press, Cambridge.
- [7] Guarino N., and Giarretta Pierdaniele, 1995, Ontologies and Knowledge Bases: Towards a Terminological Clarification, *Building and sharing very large-scale knowledge bases: Towards very large knowledge bases; Knowledge building and knowledge sharing 1995; International conference. april 1995*, Mars N., ed., Ios Pr, Amsterdam, pp. 25–32.
- [8] Gruber T. R., 1993, "A translation approach to portable ontology specifications," *Knowledge Acquisition*, **5**(2), pp. 199–220.
- [9] Borst W. N., 1997, "Construction of engineering ontologies for knowledge sharing and reuse," Dissertation, Centre for Telematics and Information Technology, University of Twente, Enschede.
- [10] Staab S., and Studer R., Eds., 2009. *Handbook on Ontologies*, Springer Berlin Heidelberg, Berlin, Heidelberg.
- [11] Horrocks I., Patel-Schneider P. F., and van Harmelen F., 2003, "From SHIQ and RDF to OWL: The making of a Web Ontology Language," *Web Semantics: Science, Services and Agents on the World Wide Web*, **1**(1), pp. 7–26.
- [12] Sprenger A., Haydn M., Ondoua S., Mosch L., and Anderl R., 2011, "Ontology-Based Information Model for the Exchange of Uncertainty in Load Carrying Structures," *Applied Mechanics and Materials*, **104**.
- [13] Zocholl M., and Anderl R., 2014, "Ontology-based Representation of Time Dependent Uncertainty Information for Parametric Product Data Models," *Proceedings of the International Conference on Knowledge Management and Information Sharing: Rome, Italy, 21-24 October 2014 ; [one of three integrated conferences that ... constitute the International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K)]*, SCITEPRESS, [S.l.], pp. 400–404.
- [14] Barbau R., Krma S. I., Fiorentini X., Rachuri S., Narayanan A. N., Foufou S., and Sriram R. D., 2010, "OntoSTEP: Enriching Product Model Data Using Ontologies," *IEEE Transactions on Knowledge and Data Engineering*, **Submitted for Publication**, **8/2010**.
- [15] Fanizzi N., d'Amato C., and Esposito F., 2008, "Statistical Learning for Inductive Query Answering on OWL Ontologies," *The Semantic Web - ISWC 2008*, Springer Berlin Heidelberg.
- [16] Schölkopf B., and Smola A. J., 2002. *Learning with Kernels: Support vector machines, regularization, optimization, and beyond*, MIT Press, Cambridge Mass. u.a.
- [17] d'Amato C., 2007, "Similarity-based Learning Methods for the Semantic Web," Dissertation, University of Bari, Bari, http://www.di.uniba.it/~cdamato/PhDThesis_dAmato.pdf.

-
- [18] Claudia d'Amato N. F. F. E., 2006, "Reasoning by analogy in description logics through instance-based learning,"
- [19] d'Amato C., Fanizzi N., and Esposito F., 2006, "A dissimilarity measure for ALC concept descriptions," *Proceedings of the 2006 ACM symposium on Applied computing*, ACM, New York, NY, p. 1695.
- [20] d'Amato C., Fanizzi N., and Esposito F., Eds., 2008. *Query Answering and Ontology Population: An Inductive Approach: The Semantic Web: Research and Applications*, Springer Berlin Heidelberg.
- [21] Zied Bouraoui, Shoaib Jameel, and Steven Schockaert, "Inductive Reasoning about Ontologies Using Conceptual Spaces,"
- [22] Gärdenfors P., 2000. *Conceptual spaces: the geometry of thought*, MIT Press, Cambridge, Massachusetts.
- [23] Jameel S., and Schockaert S., 2006, Entity Embeddings with Conceptual Subspaces as a Basis for Plausible Reasoning, *Proceedings of the 22nd Conference on Artificial Intelligence*.
- [24] Mascardi, V. ,Cordi, V., Rosso, P., 2007, "A Comparison of Upper Ontologies," <https://pdfs.semanticscholar.org/4f28/6fdf9280449588b9d3781c9c897da28e0cff.pdf>.
- [25] Niles I., and Pease A., 2001, Towards a Standard Upper Ontology, *Formal ontology in information systems: Collected papers from the second International Conference, October 17th-19th, 2001, the Cliff House, Ogunquit, Maine, USA*, Welty C., and Smith B., eds., Association for Computing Machinery, New York, pp. 2–9.
- [26] Pease A., 2011. *Ontology: A practical guide*, Articulate Software Press, Angwin, CA.
- [27] Derrac J., and Schockaert S., 2015, "Inducing semantic relations from conceptual spaces: A data-driven approach to plausible reasoning," *Artificial Intelligence*, **228**, pp. 66–94.
- [28] Minervini P., d'Amato C., Fanizzi N., and Tresp V., 2016, "Discovering Similarity and Dissimilarity Relations for Knowledge Propagation in Web Ontologies," *Journal on Data Semantics*, **5**(4), pp. 229–248.