PUCRS

ESCOLA POLITÉCNICA
PROGRAMA DE PÓS-GRADUAÇÃO EM CIÊNCIA DA COMPUTAÇÃO
DOUTORADO EM CIÊNCIA DA COMPUTAÇÃO

DANIELA SCHMIDT

ALIGNING TOP-LEVEL AND DOMAIN ONTOLOGIES

Porto Alegre
2020

PÓS-GRADUAÇÃO - *STRICTO SENSU*

AD VERVM DVCIT

Pontifícia Universidade Católica
do Rio Grande do Sul

PONTIFICAL CATHOLIC UNIVERSITY OF RIO GRANDE DO SUL
SCHOOL OF TECHNOLOGY
COMPUTER SCIENCE GRADUATE PROGRAM

ALIGNING TOP-LEVEL AND DOMAIN ONTOLOGIES

DANIELA SCHMIDT

Dissertation submitted to the Pontifical Catholic University of Rio Grande do Sul in partial fulfillment of the requirements for the degree of Ph. D. in Computer Science.

Advisor: Prof. Dra. Renata Vieira
Co-Advisor: Prof. Dra. Cássia Trojahn dos Santos

Ficha Catalográfica

S349a Schmidt, Daniela

Aligning Top-level and Domain Ontologies / Daniela Schmidt. – 2020.
91.
Tese (Doutorado) – Programa de Pós-Graduação em Ciência da Computação, PUCRS.

Orientadora: Profa. Dra. Renata Vieira.
Co-orientadora: Profa. Dra. Cássia Trojahn dos Santos.

1. Ontology matching. 2. top-level ontology. 3. domain ontology. I. Vieira, Renata. II. dos Santos, Cássia Trojahn. III. Título.

Elaborada pelo Sistema de Geração Automática de Ficha Catalográfica da PUCRS
com os dados fornecidos pelo(a) autor(a).
Bibliotecária responsável: Clarissa Jesinska Selbach CRB-10/2051

Daniela Schmidt

Aligning top-level and domain ontologies

This Dissertation has been submitted in partial fulfillment of the requirements for the degree of Doctor of Computer Science, of the Graduate Program in Computer Science, School of Technology of the Pontifical Catholic University of Rio Grande do Sul.

Sanctioned on January 23th, 2020.

COMMITTEE MEMBERS:

Prof. Dra. Mara Abel (PPGC/UFRGS)

Prof. Dr. Sandro Rigo (PPGCA/Unisinos)

Prof. Dra. Sabrina dos Santos Marczak (PPGCC/PUCRS)

Prof. Dra. Cássia Trojahn dos Santos (IRIT-FR – Co-Advisor)

Prof. Dra. Renata Vieira (PPGCC/PUCRS - Advisor)

Dedico este trabalho a meus pais Nelson e Helga.

ACKNOWLEDGMENTS

Agradeço primeiramente aos meus pais Nelson e Helga que são uma fonte imensurável de apoio e incentivo extremamente importante na minha caminhada. Pela compreensão dispendida nos momentos de ausência dedicados para a realização deste trabalho.

Agradeço os meus irmãos e familiares pelo apoio incondicional e ao Edson pelo companheirismo.

Agradeço imensamente as minhas queridas orientadoras, professoras Renata Vieira e Cassia Trojahn dos Santos, pela confiança depositada no meu trabalho e pelo apoio, ensinamentos e amizade.

Aos meus colegas do laboratório de Processamento de Linguagem Natural (PLN), em especial aos amigos Evandro, Sandra, Daniela que tive o privilégio de conviver ativamente durante os anos de andamento do doutorado.

Aos bolsistas de iniciação científica que participaram ativamente desse projeto, em especial ao Rafael e Henrique.

Agradeço a Pontifícia Universidade Católica do Rio Grande do Sul (PUCRS), por ter me acolhido nesses sete anos de experiência na Pós-graduação. Aos professores e funcionários que tive o privilégio de conhecer e trocar conhecimentos.

Por fim, a Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (Capes) pela oportunidade de cursar o Doutorado.

ALINHANDO ONTOLOGIAS DE TOPO E DE DOMINIO

RESUMO

Ontologias têm sido aplicadas em diversas áreas, motivado principalmente pela necessidade de representar, compartilhar e reusar conhecimento. Ontologias de topo têm um papel fundamental na construção e integração de ontologias de domínio, fornecendo um modelo de referência fundamentado que pode ser compartilhado por diferentes domínios. Os esforços na área de alinhamento de ontologias se concentram essencialmente no alinhamento de ontologias de domínio. No entanto, o problema de alinhar ontologias de domínio e topo é abordado em menor proporção, principalmente devido aos diferentes níveis de abstração destes tipos de ontologias. Assim, este trabalho se concentra em preencher uma lacuna nessa área. Uma abordagem automática para alinhar ontologias de topo e de domínio é proposta. A abordagem é baseada em alinhamentos existentes entre recursos lexicais e ontologias de topo, sendo usados como uma camada intermediária. As técnicas utilizadas no nosso processo de alinhamento baseiam-se em similaridade de string (Lesk e Leveinstein) e no uso de modelos pré-treinados de word embedding (Glove e GoogleNews). A análise de resultados foi avaliada comparando as técnicas propostas no trabalho, e ferramentas desenvolvidas para alinhamento automático de ontologias. Embora as ferramentas utilizadas na comparação não foram desenvolvidas com o foco em alinhamento de ontologias de topo e domínio, são ferramentas estado da arte na área. Os resultados foram avaliados considerando alinhamentos de referência.

Palavras-Chave: Alinhamento de Ontologias, Ontologia de topo, Ontologia de domínio.

ALIGNING TOP-LEVEL AND DOMAIN ONTOLOGIES

ABSTRACT

Ontologies have been applied in diverse areas, mainly motivated by the need to represent, to share and to reuse knowledge. Top-level ontologies play a fundamental role in the construction and integration of domain ontologies, providing a well-founded reference model that can be shared across knowledge domains. Efforts in ontology matching field have been particularly dedicated to domain ontologies, however the problem of matching domain and top-level ontologies has been addressed to a lesser extent, particularly due to their different levels of abstraction. Therefore, this work aims at filling the gap in this area. An automatic approach to align top and domain ontologies is proposed. Our approach is based on existing alignments between lexical resources and top ontologies, being used as an intermediate layer. The results was evaluated by comparing the techniques proposed in our work regarding tools developed for automatic alignment of ontologies. Although the adopted tools were not developed with a focus on aligning top and domain ontologies, they are the state of art tools in these field. The results were evaluated considering reference alignments.

Keywords: Ontology matching, top-level ontology, domain ontology.

LIST OF FIGURES

Figure 2.1 – Main concepts of DOLCE. 32
Figure 2.2 – Main concepts of SUMO 33
Figure 2.3 – The matching process according to Euzenat and Shvaiko [5]. 35
Figure 2.4 – The matching techniques classification according to Euzenat and Shvaiko [5]. 36
Figure 2.5 – Example of WordNet senses from Conference 40
Figure 2.6 – Simplified Lesk algorithm [28]. 41
Figure 2.7 – Levenstein distance example [28]. 41
Figure 4.1 – General view of the proposed approach. 51
Figure 4.2 – Example of WordNet synsets linked to DOLCE. 55

LIST OF TABLES

Table 3.1 – Summary of the approaches on chronological order 47
Table 5.1 – Example of spreadsheet adopted by the evaluators. 61
Table 5.2 – Conflicts solved after discussion between annotators. Boldface concepts represent the conflicts and final results are indicated with a underline. 62
Table 5.3 – Excerpt of DUL correspondences. Boldface concepts represent the chosen correspondence. 65
Table 5.4 – SUMO to CORA alignments. 66
Table 5.5 – DUL to SOSA and SSN alignments. 67
Table 6.1 – Summary of domain ontologies. 70
Table 6.2 – Word sense disambiguation techniques, pair of ontologies and found correspondences considering WordNet previous alignments. 72
Table 6.3 – Word sense disambiguation techniques, pair of ontologies and found correspondences without previous WN alignments. 72
Table 6.4 – Tasks and found correct correspondences by each disambiguation method 73
Table 6.5 – Precision, Recall, and F-measure for each disambiguation method. . . 74
Table 6.6 – Precision, Recall, and F-measure for each disambiguation method, compiled by domain. 74
Table 6.7 – Excerpt of correspondences found by our approach considering the similarity of Word Embedding Models. 75
Table 6.8 – Summary of correspondences found by OAEI matching tools. 76
Table 6.9 – Precision, Recall, and F-measure for each disambiguation method, comparing with OAEI tools. 76
Table 6.10 – Examples of Found Correspondences. 78

LIST OF ACRONYMS

RDF – Resource Description Framework

OWL – Ontology Web Language

OAEI – Ontology Alignment Evaluation Initiative

DOLCE – Descriptive Ontology for Linguistic and Cognitive Engineering

DL – DOLCE-Lite

DLP – DOLCE-Lite-Plus

DUL – DOLCE-Ultralite

LOA – Laboratory for Applied Ontology

KIF – Knowledge Interchange Format

BFO – Basic Formal Ontology

IFOMIS – Institute for Formal Ontology and Medical Information Science

SNAP – A series of snapshot ontologies

SPAN – A single videoscopic ontology

GFO – General Formal Ontology

SUMO – Suggested Upper Merged Ontology

TP – True positive

FP – False positive

TN – True Negative

FN – False Negative

CONTENTS

1 INTRODUCTION 25
1.1 RESEARCH HYPOTHESIS 26
1.2 RESEARCH GOALS 26
1.3 OUTLINE 27

2 BACKGROUND 29
2.1 ONTOLOGY 29
2.1.1 TOP-LEVEL ONTOLOGIES 30
2.2 ONTOLOGY MATCHING 35
2.2.1 MATCHING EVALUATION 38
2.3 SEMANTIC RESOURCES AND SENSE DISAMBIGUATION 39
2.3.1 WORDNET 39
2.3.2 WORD SENSE DISAMBIGUATION 40
2.4 WORD EMBEDDING 42
2.5 CONCLUSIONS 42

3 RELATED WORK 43
3.1 MATCHING TOP ONTOLOGIES 43
3.2 MATCHING TOP ONTOLOGIES TO LEXICONS 44
3.3 MATCHING DOMAIN ONTOLOGIES VIA TOP ONTOLOGIES 45
3.4 MATCHING DOMAIN ONTOLOGIES TO TOP ONTOLOGIES 46
3.5 DISCUSSION 47
3.6 CONCLUSIONS 49

4 PROPOSED APPROACH 51
4.1 PRE-PROCESSING 52
4.2 WORD SENSE DISAMBIGUATION 52
4.2.1 LESK MEASURE (EXACT MATCH) 53
4.2.2 LEVENSHTEIN DISTANCE (STRING DISTANCE) 53
4.2.3 WORD EMBEDDING 54
4.3 ALIGNMENT METHOD - VIA PREVIOUS WORDNET ALIGNMENTS 55
4.3.1 DOLCE CORRESPONDENCE IDENTIFICATION 55
4.3.2 SUMO CORRESPONDENCE IDENTIFICATION 56

4.4 ALIGNMENT METHOD - WITHOUT PREVIOUS ALIGNMENTS 57
4.5 TECHNICAL ASPECTS 57
4.6 CONCLUSIONS 58

5 REFERENCE ALIGNMENTS 59
5.1 ALIGNING CONFERENCE ONTOLOGIES WITH SUMO 59
5.1.1 SUMO ALIGNMENT METHODOLOGY 60
5.1.2 INDIVIDUAL GENERATION OF CORRESPONDENCES 61
5.1.3 VALIDATING THE CORRESPONDENCES 61
5.1.4 DISCUSSION 62
5.2 ALIGNING CONFERENCE ONTOLOGIES WITH DOLCE 64
5.3 EXISTING ALIGNED ONTOLOGIES 64
5.3.1 CORA ONTOLOGY 65
5.3.2 SSN (W3C SEMANTIC SENSOR NETWORK ONTOLOGY) 65
5.4 CONSLUSIONS 67

6 EXPERIMENTS 69
6.1 MATERIAL AND METHODS 69
6.1.1 ONTOLOGIES 69
6.1.2 EXTERNAL RESOURCES 70
6.1.3 WORD EMBEDDING MODELS 71
6.1.4 OAEI TOOLS 71
6.2 EXPERIMENTS DESIGN 71
6.2.1 APPROACHES CONSIDERING PREVIOUS WORDNET ALIGNMENTS 72
6.2.2 APPROACHES WITHOUT PREVIOUS WORDNET ALIGNMENTS 72
6.3 RESULTS 72
6.3.1 RESULTS CONSIDERING PREVIOUS WORDNET ALIGNMENTS 73
6.3.2 RESULTS WITHOUT PREVIOUS WORDNET ALIGNMENTS 74
6.3.3 OAEI 2017 MATCHING TOOLS 75
6.4 DISCUSSION AND ERROR ANALYSIS 76
6.5 CONCLUSIONS 79

7 CONCLUSION 81
7.1 FUTURE WORK 82
7.2 PUBLICATIONS 82

REFERENCES 85

1. INTRODUCTION

Ontologies have been applied in diverse areas motivated mainly by the need to represent, share and reuse knowledge. Guarino [16] classifies ontologies according to their “level of generality": *(i) top-level ontologies* describe very general concepts (e.g., space, time, object, etc.), which are independent of a particular problem or domain, these are also named upper or foundational ontologies [59, 21]; *(ii) domain ontologies* and *task ontologies* describe, respectively, the entities and other information related to a generic domain (e.g., biology or aeronautic), or a generic task or activity (e.g., diagnosis), ideally by specializing the concepts represented in top-level ontologies; and finally *(iii) application ontologies*, which describe roles played by entities when performing an activity (which are, respectively, described by domain and activity ontologies).

Ontologies can be modeled with different models and they can be expressed in various kinds of languages. At the beginning of the 1990s, ontologies were built using techniques based on frames and first-order logic. In the last years, other knowledge representation techniques based on description logics have been used to build ontologies and description logic languages like OIL, DAML+OIL and OWL have appeared in the context of the Semantic Web [11]. Ideally, ontologies developed on the basis of these languages should be complemented by a methodology based on top-level ontologies [20], since they are usually equipped with a rich axiomatic layer. Hence, while the rich semantics and formalization of top-level ontologies are important requirements for ontology design [41], they act as well as semantic bridges supporting very broad semantic interoperability between ontologies [38, 39]. In that sense, they play a key role in ontology matching, which is the process of finding correspondences between entities from different ontologies. The advantage of top-level ontologies is to provide a well-founded reference model that can be shared across knowledge domains. Moreover, the clarity in semantics and the rich formalization of these ontologies are fundamental requirements for ontology development [30] improving ontology quality.

However, most efforts in ontology matching have been particularly dedicated to domain ontologies and the problem of matching domain and top-level ontologies has been addressed to a lesser extent. This problem poses different challenges in this field, particularly due to the different levels of abstraction of these ontologies. This is a complex task, that requires knowledge about the semantic context of concepts, which goes beyond the frontiers of what is encoded in the ontology. In particular, the differences in the abstraction levels of domain and top-level ontologies require a change of focus from finding equivalence relation to the identification of subsumption relations. In fact, when having different levels of abstraction it might be the case that the matching process should focus in finding subsumption rather than equivalence correspondences, since the top-level ontology has concepts at a higher level. This, so far, is largely neglect by most matching systems. Approaches dealing with this task are mostly based on manual matching [1, 41].

In order to evaluate the quality and correctness of the generated alignments in the process of top-level and domain ontology matching, reference alignments (also called gold standard) are required. Reference alignments could be developed manually by experts or in a semi-automatic way, where the resulting alignment from matching systems is evaluated against a manual analysis. However, in the current state of the art, the available reference alignments frequently involve just domain ontologies such as the OAEI[1] reference alignments.

The task of align top-level and domain ontologies manually is a hard work task, that deal with the knowledge and interpretation of experts. In this way, a tool that helps automatizing the task is an improvement to the ontology alignment area. Considering the discussion above, this thesis aims to contribute with the problem of matching domain and top-level ontologies. For that aim we propose an approach to automatically align domain and top-level ontologies. More specifically, we intend to contribute with *(i)* the creation of a reference alignment for domain and top-level ontologies that might be integrated in future evaluation campaigns; *(ii)* an approach to match top-level and domain ontologies automatically; and *(iii)* a prototype implementing our matching approach. In this way, we are reducing the effort to manually align ontologies, once the matcher generates the correspondences and reference alignments could be adopted to evaluate the generated alignments.

1.1 Research Hypothesis

Considering the depicted discussed challenges of matching top-level and domain ontologies, we define a research hypothesis which will be a guide of our research work as follow: the existing alignments between lexical database WordNet and top-level ontologies could be a way to improve and obtain gains in relation to the state-of-the-art matching systems in the task of matching domain and top-level ontologies.

1.2 Research Goals

Goal 1. Propose an approach to match domain and top-level ontologies based on existing alignments between top-level ontologies and WordNet. To achieve this goal, we intend to:

- Study previous alignments between top-level ontologies and WordNet;
- Investigate techniques to generate concept relationships to incorporate into our approach.

[1] http://oaei.ontologymatching.org/

Goal 2. Develop a prototype to evaluate our proposed model. To achieve this goal, we intend to:
- Implement a matching system based on the issues investigated on goal 1;

Goal 3. Evaluate the proposed prototype with a reference alignment. To achieve this goal, we intend to:
- Create a reference alingment;
- Evaluate the performance of the approach in terms of precision, recall, and F-measure;
- Compare our prototype with available matching tools.

1.3 Outline

This thesis is structured as follows: Chapter 2 introduces the theoretical background on ontology and ontology matching. Then, we discuss about available state-of-the-art matching systems and their evaluation. Finally, we present the semantic resources adopted in our work.

Chapter 3 discusses related work. In this chapter, we intend to present initiatives to match domain ontologies with top-level ontologies. Our goal is to identify how the alignments are created and investigate the use of external sources in the ontology matching task. We conclude the chapter with a summary of related works and explain the gaps that motivated our work.

Chapter 4 presents our proposed approach. We start describing initial processing phases, then, we describe adopted word sense disambiguation techniques including in our algorithm. Next, our alignment method is described, and finally, technical aspects of the implementation are detailed.

Chapter 5 describes the methodology followed for constructing the reference alignments.

Chapter 6 details our evaluation process. Our obtained results are also presented and discussed in the chapter.

Chapter 7 ends the thesis presenting the conclusions and perspectives for future work.

2. BACKGROUND

This chapter is divided in three sections. The first section introduces the background on ontology. The second section presents the main concepts of ontology matching, briefly describing classical matching evaluation metrics. Finally, the third section introduces the semantic resources that have been adopted in our matching approach.

2.1 Ontology

Gómez-Pérez et al. [11] start the discussion about ontologies with some questions such as: What are things? What is the essence that remains inside things even when they change (changes in their color, changes in their size, etc.)? Do concepts (book, tree, table, etc.) exist outside our mind? How can the entities of the world be classified? According to the authors, these are some of the questions that Ontology, the philosophy of existence, has tried to answer for thousands of years. Hence, the term ontology was originally the philosophical study of reality to define which things exist and what we can say about them. The term ontology was also adopted in other research areas, including Computer Science where it is defined as an "explicit specification of a conceptualisation" [14]. A conceptualisation stands for an abstract model of some aspect of the world where properties of important concepts and relationships are specified.

Ontology in Computer Science is used to capture and express features of things in a computational understandable way. It can be developed using different techniques and languages, furthermore, as more the essence of things is captured, more possible it is for the ontology to be shared. So, the essence of the Greek philosophy, which tried to make unity compatible with variety is an important point to take into account, although in a different context [11].

According to [67] *"an ontology is a hierarchically structured set of terms for describing a domain that can be used as a skeletal foundation for a knowledge base"*, hence, the same ontology can be used for building many knowledge bases that share the same taxonomy. Moreover, it is possible to add low level terms as sub-concepts or upper level terms as super-concept to cover new areas.

Ontologies can be distinguished as *lightweight* and *heavyweight*. *Lightweight* ontologies can be seen as taxonomies that include concepts and their taxonomies, relationships between concepts, and properties that describe concepts. *Heavyweight* ontologies add axioms and constraints to *lightweight* ontologies [11].

According to Guarino [16], ontologies can also be classified according to their "level of generality", in top-level, domain, task, and application ontologies. In this work, the distinction between domain and top ontologies is an important one and it is discussed in the sequence.

2.1.1 Top-level Ontologies

A top-level ontology is a high-level and domain independent ontology. It describes general concepts (e.g., physical object) and relations (e.g., parthood), which are independent of a particular domain. The concepts expressed are intended to be basic and universal to ensure generality and expressiveness for a wide range of domains. It is often characterized as representing common sense concepts and it is limited to concepts which are meta, generic, abstract and philosophical. There are two approaches for the use of top-level ontologies [59], *top-down* and *bottom-up*. Top-down approach uses the ontology as a foundation for deriving concepts in the domain ontology. In this way, we take the advantage of the knowledge and experience already expressed in the top-level ontology. In a bottom-up approach, one usually matches a new or existing domain ontology to a given top-level ontology. This approach represents more challenges since inconsistencies may exist between domain and top-level ontologies [59].

According to Gómez-Pérez et al. [11] some important philosophical notions could be considered in the top-level ontologies, such as rigidity, identity, dependency, and unity which are explained in the context of the OntoClean method. The OntoClean [17] is defined as *"a methodology for validating the ontological adequacy of taxonomic relationships based on highly general ontological notions drawn from philosophy, like essence, identity, and unity, which are used to characterize relevant aspects of the intended meaning of the properties, classes, and relations that make up an ontology."* Some of these notions are explained below:

- Essence and Rigidity [17, 11]: a property of an entity is considered essential to that entity if it must be true of it in every possible world, for instance, every human must have a brain in every possible world. The rigidity is defined according to the idea of essence. Hence, a property is rigid (+R) if it is essential to all its possible instances. An instance of a rigid property cannot stop being an instance of that property in a different world. For example, having a brain may be essential to humans. If a property is not essential for all their instances, it is classified as anti-rigid (R̃), while non-rigid (-R) properties could acquire or lose their instances depending on the state of affairs at hand.

- Identity and Unity [17, 11]: a property carries an identity criteria (+I) if and only if all its instances can be (re)identified by means of a suitable sameness relation. It refers to the problem of being able to recognize individual entities in the world as being the same and are conditions used to determine equality (sufficient conditions) and that are entailed by equality (necessary conditions). On the other hand, unity (+U) refers to being able to recognize all the parts that form an individual entity. It refers to the problem of describing the parts and boundaries of objects, such that we know in general what is part of the object, what is not, and under what conditions the object is *whole*. For instance, *being an ocean* is a property that picks up whole objects, as its instances, such as *the Atlantic Ocean*. On the other hand, *being (an amount of) water* does not have wholes as instances, hence, knowing an entity is an amount of water does not tell us anything about its parts and it is classified as anti-unity (Ũ).

- Dependency [11]: a property P is constantly dependent (+D) if and only if for all its instances, there are something on which the instances are constantly dependent. Hence, an individual x is constantly dependent on the individual y if and only if, at any time, x cannot be present unless y is fully present, and y is not part of x. For example, a *hole* in a *wall* is constantly dependent on the *wall*.

Top-level ontologies are usually modeled based on such kinds of philosophical notions. In the literature it is possible to identify many efforts to develop top-level ontologies. Guarino [15] proposed a set of principles for an unified top-level ontology based on clarity, semantic rigour, generality, and commonsense and then, shows a preliminary proposal for top-level concepts based on these principles.

Since then, several top-level ontologies have been proposed. Mascardi et al. [38] presents an analyse of 7 top-level ontologies called BFO, Cyc, DOLCE, GFO, PROTON, Sowa's ontology, and SUMO. The authors analysis include criteria as: dimension, implementation language, modularity, developed applications licensing among others. A summary of them is presented bellow:

- DOLCE [8], Descriptive Ontology for Linguistic and Cognitive Engineering, has been proposed by Nicola Guarino and his team at LOA (Laboratory for Applied Ontology). DOLCE is the first module of the WonderWeb Foundational Ontologies Library. The focus of DOLCE is to grasp the underlying categories of human cognitive tasks and the socio-cultural environment. It is an ontology of particulars and include concepts such as abstract quality, abstract region, physical object, process, and so on.

  DOLCE was designed to include the most reusable and widely applicable upper-level categories, it is rigorous in terms of axiomatization and extensively researched and documented. It is an ontology of particulars which has four top-level concepts: *endurant*, *perdurant*, *quality*, and *abstract*. *Endurants* represent objects or substances while *perdurants* corresponds to events or processes. The main relation between *endurants* and *perdurants* is that of participation, e.g., a person which is an *endurant*, may participate in a discussion, which is a *perdurant*. *Qualities* can be seen as the basic entities that we can perceive or measure, e.g., shapes, colors, sizes, etc. *Abstracts* do not have spatial or temporal qualities, and they are not qualities themselves.

DOLCE is implemented in First Order Logic and Knowledge Interchange Format (KIF). It has been adapted and re-engineered according to the expressivity of OWL. Their efforts resulted in various OWL models, among which DOLCE Lite. DOLCE is composed by around one hundred concepts. DOLCE is freely available at *http://www.loa.istc.cnr.it/ old/DOLCE.html*. The ontology has many variations, such as DOLCE-Lite [9], which is an OWL-DL fragment of DOLCE, this fragment is composed by 37 concepts and 70 object properties. DOLCE-Lite has been extended in modules for representing information, communication, plans, and with some domain information for representing e.g. legal, biomedical notions. The combination of DOLCE-Lite (DL) and the mentioned additional modules is called DOLCE-Lite-Plus (DLP), and it is available in *http://www.loa.istc.cnr.it/ old/ontologies/DLP_397.owl*. DOLCE+DnS Ultralite ontology (DUL) is a simplification and extension of DOLCE and Descriptions and Situations ontologies created by Aldo Gangemi [7]. Figure 2.1.1 shown the main concepts of DOLCE ontology.

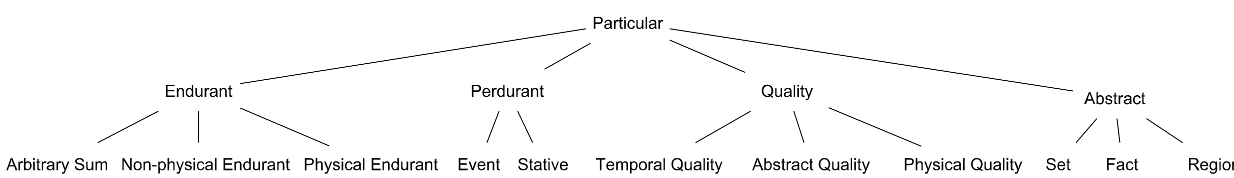

Figure 2.1 – Main concepts of DOLCE.

- SUMO [45]: Suggested Upper Merged Ontology is an upper level ontology that has been proposed as a starter document for The Standard Upper Ontology Working Group, an IEEE working group of collaborators from the fields of engineering, philosophy, and information science. SUMO provides definitions for general-purpose terms and acts as a foundation for more specific domain ontologies. It is being used for research and applications in search, linguistics and reasoning. SUMO is developed in first order logic language and OWL. It is an ontology of particulars and universals which has two top-level concepts: *physical* and *abstract*. *Physical* represent an entity that has a location in space-time. An *abstract* can be said to exist in the same sense as mathematical objects such as sets and relations, but they cannot exist at a particular place and time without some physical encoding or embodiment. The OWL version is composed by 4.558 concepts and 778 properties. The ontology is freely available at http://www.adampease.org/OP. Figure 2.1.1 shown the main concepts of SUMO ontology.

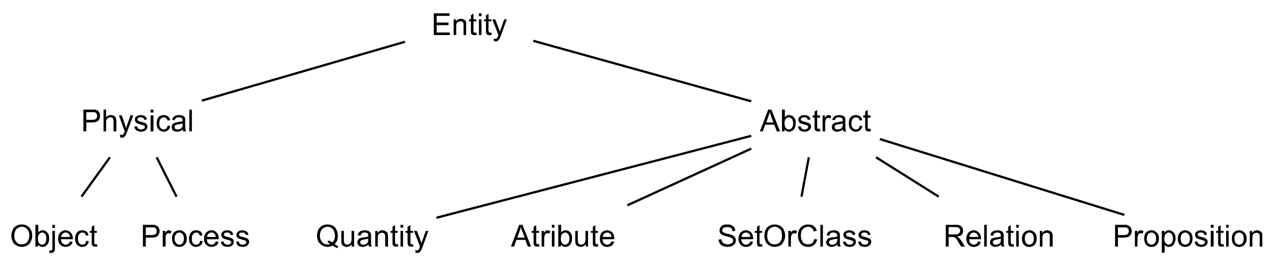

Figure 2.2 – Main concepts of SUMO

Our work explores DOLCE [7] and SUMO [45]. Our choice is motivated by the following: (i) these ontologies are widely adopted in research projects; (ii) there are extensive documentations on them; (iii) both are previously aligned with WordNet. A comparison of DOLCE and SUMO is presented in [50] and discussed in [38]. According to the authors, DOLCE has a similar purpose and business process to SUMO in that it is a free research project for use in both natural language tasks and inference. DOLCE is an "ontology of particulars"; it does have universals (classes and properties), but the claim is that they are only employed in the service of describing particulars. In contrast, SUMO could be described as an ontology of both particulars and universals. It has a hierarchy of properties as well as classes. This is a very important feature for practical knowledge engineering, as it allows common features like transitivity to be applied to a set of properties, with an axiom that is written once and inherited by those properties, rather than having to be rewritten, specific to each property. Other differences include DOLCE's use of a set of meta-properties as a guiding methodology, as opposed to SUMO's use and formal definition of such meta-properties directly in the ontology itself. With respect to SUMO, DOLCE does not include such items as a hierarchy of process types, physical objects, organisms, units and measures,and event roles.

Next, we summarize the main characteristics of other top-level ontologies, following the work of Mascardi et al. [38] and including one more foundational ontology, UFO, proposed by Guizzardi [19]. They are not in the scope of this thesis for the purpose of matching analysis, but it is important to acknowledge their role in this area.

- BFO [13]: Basic Formal Ontology has been proposed by Barry Smith and Pierre Grenon in their team at IFOMIS (Institute for Formal Ontology and Medical Information Science). BFO is designed for use in supporting information retrieval, analysis and integration in scientific and other domains. BFO consists in two sub-ontologies: *(i)* SNAP (a series of snapshot ontologies (O_{t_i})), indexed by time and *(ii)* SPAN (a single videoscopic ontology (O_v)). An O_{t_i} is an inventory of all entities existing at a time, while an O_v is an inventory of all processes unfolding through time. BFO contains 1 top connecting class ("Entity"), 18 SNAP classes, and 17 SPAN classes for a total of

36 classes which are, in version 1.0 of its implementation, connected via the is_a relation. The ontology is implemented in the OWL language and it is freely available at http://ifomis.uni-saarland.de/bfo.

- Cyc [18]: The Cyc ontology is developed by Cycorp company. The Cyc Knowledge Base (KB) is a formalised representation of facts, rules of thumb, and heuristics for reasoning about the objects and events of everyday life. The KB consists of terms and assertions which relate those terms. These assertions include both simple ground assertions and rules. The Cyc KB is divided into thousands of "microtheories" focused on a particular domains of knowledge, a particular level of detail, a particular interval in time, etc. The Cyc Knowledge Base contains more than 300,000 concepts and nearly 3,000,000 assertions (facts and rules), using more than 15,000 relations. Cyc is implemented in the CycL language, but it is possible to export to OWL format. Cyc is a commercial product available at http://www.cyc.com. There is an open source release called OpenCyc available at http://sw.opencyc.org.

- GFO [25]: General Formal Ontology is a top-level ontology for conceptual modeling that has been proposed by the Onto-Med Research Group. It includes elaborations of categories such as objects, processes, time and space, properties, relations, roles, functions, facts, and situations. The work is in progress on the integration with the notion of levels of reality in order to more appropriately capture entities in the material, mental, and social areas. GFO is implemented in KIF and OWL languages. The OWL version of GFO consists of 79 classes and 67 properties. GFO is freely available at http://www.onto-med.de/ontologies/gfo/index.jsp.

- PROTON [69]: PROTo ONtology was developed in the SEKT project as a lightweight upper-level ontology, serving as a modeling basis for tasks in different domains including semantic annotation, indexing, and retrieval of documents. PROTON contains about 300 classes and 100 properties. PROTON is implemented in OWL language and it is freely available at http://ontotext.com/proton.

- SOWA's Ontology [65]: The ontology was developed by J.F.Sowa. The basic categories were derived from a variety of sources in logic, linguistics, philosophy, and artificial intelligence. Sowa's ontology is not based on a fixed hierarchy of categories, but on a framework of distinctions, from which the hierarchy is generated automatically. Sowa's ontology is implemented in first order modal language and KIF. The KIF version contains about 30 classes and 5 properties. The KIF version is freely available at http://www.jfsowa.com/ontology.

- UFO [19, 21]: Unified Foundational Ontology was developed by Giancarlo Guizzardi. The ontology development started as a unification of the GFO (Generalized Formalized Ontology) and the top-Level ontology of universals underlying OntoClean. UFO comprises areas as cognitive science, linguistics and philosophical logics. Nowadays, the ontology is divided in three parts representing different aspects of reality. UFO-A is an ontology of endurants dealing with aspects of structural conceptual modeling comprising theories of types and taxonomic structures, part-whole relations, among others. UFO-B is an ontology of perdurants (events and processes) dealing with aspecs as perdurant mereology, temporal ordering of perdurants, object participation in perdurants, among others. And UFO-C is an ontology of intentional and social entities, which is constructed on top of UFO-A and UFO-B, and which addresses notions such as beliefs, desires, intentions, goals, actions, commitments and claims, social roles and social particularized relational complexes (Social Relators), among others. UFO is a work in progress and it is possible to find many application in the literature [22].

2.2 Ontology Matching

According to Euzenat and Shvaiko [5], an alignment is defined as a set of correspondences between two ontologies. The alignment is the output of the matching process.

The matching operation takes as input two ontologies o_s (source) and o_t (target) and an (possibly empty) alignment A to be completed. The output is an alignment A' which express a set of correspondences belonging to source and target ontologies. Euzenat and Shvaiko [5] given a definition of matching process as:

Definition 1 (Matching process) *The matching process can be seen as a function f which, from a pair of ontologies to match o and o', an input alignment A, a set of parameters p and a set of oracles and resources r, returns an alignment A' between these ontologies:*

$$A' = f(o, o', A, p, r)$$

The matching process is illustrated in the Figure 2.3.

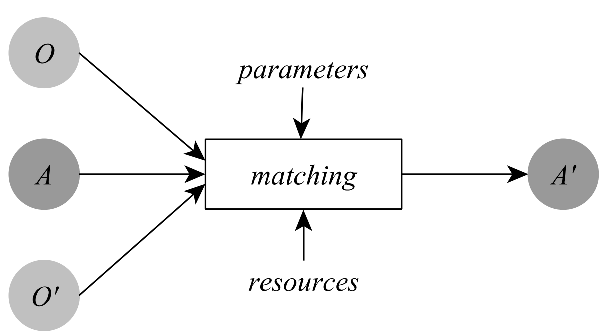

Figure 2.3 – The matching process according to Euzenat and Shvaiko [5].

A definition of correspondence is given by Euzenat and Shvaiko [5], as follows:

Definition 2 (Correspondence) *A correspondence can be defined as $\langle e_s, e_t, r, n \rangle$, such that: e_s and e_t are entities (e.g., concepts, properties, instances) of o_s and o_t, respectively; r is a relation holding between two entities e_s and e_t, (for instance,* equivalence*,* subsumption*,* disjointness*,* overlapping*); and n is a confidence measure number in the [0;1] range.*

Different matching approaches have emerged from the literature [29, 5]. While *terminological* methods lexically compare strings (tokens or n-grams) used in naming entities (or in the labels and comments of entities), *semantic* methods utilize model-theoretic semantics to determine whether or not a correspondence exists between entities. Approaches may consider the *internal* ontological structure, such as the range of their properties, their cardinality, and their transitivity and/or symmetry features, or alternatively the *external* ontological structure, such as the position of entities within the ontological hierarchy. The instances of concepts can also be compared using statistical, probabilistic and linguistic approaches. Moreover, ontology matching systems rely on several approaches.

Euzenat and Shvaiko [5] present a classification of matching techniques as we can see in Figure 2.4. The upper classification is based on granularity and input interpretation, the lower classification is based on the kind of input, and the middle layer features classes of basic techniques.

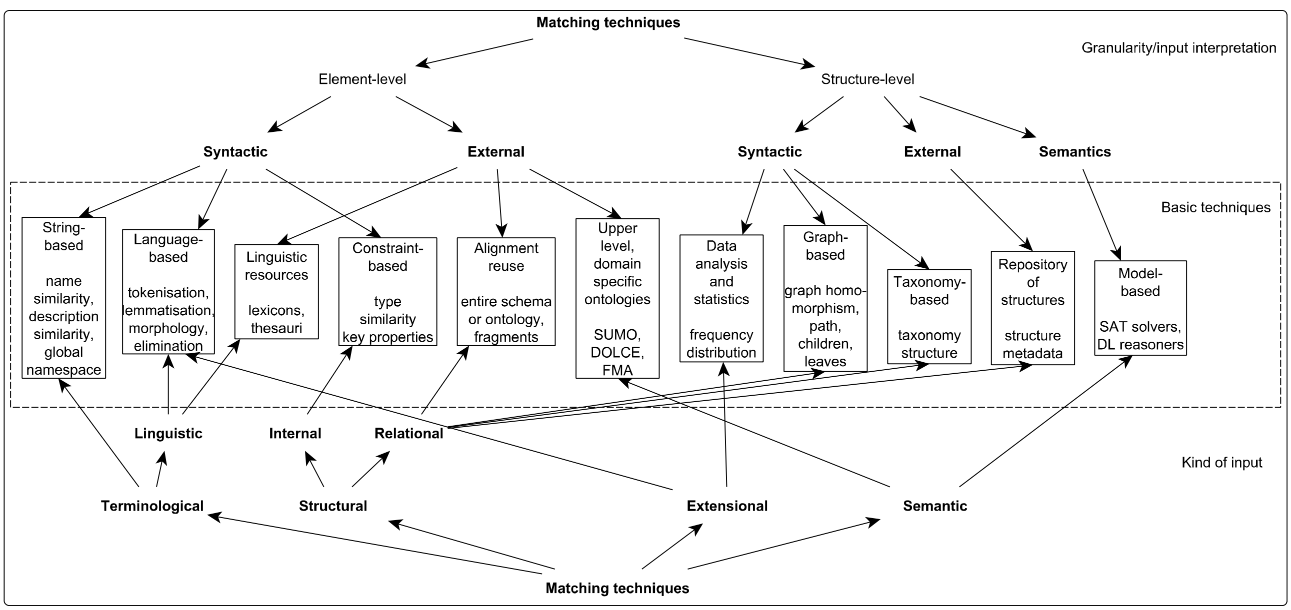

Figure 2.4 – The matching techniques classification according to Euzenat and Shvaiko [5].

Element-level techniques compute correspondences by analysing entities of those entities in isolation, ignoring their relations with other entities. Structure-level techniques compute correspondences by analysing how entities or their instances appear together in a structure. Syntactic techniques usually lexically compare strings (tokens or n-grams) used in naming entities (or in the labels and comments concerning entities). External are the techniques exploiting auxiliary (external) resources of a domain and common knowledge in order to interpret the input. Semantic techniques use some formal semantics to interpret the input and utilise model-theoretic semantics to determine whether or not a correspondence exists between two entities.

Element-level techniques include more specific techniques such as:

String-based techniques are used to match names and annotations of ontology entities. They are typically based on the following intuition: the more similar the strings, the more likely they are to denote the same concepts. Usually, distance functions map a pair of strings to a real number, where a smaller value of the real number indicates a greater similarity between the strings.

Language-based techniques consider names as words in some natural language. They are based on natural language processing techniques exploiting morphological properties of the input words. Usually, they are applied to names of entities before running string-based or lexicon-based techniques in order to improve their results.

Constraint-based techniques are algorithms which deal with the internal constraints being applied to the definitions of entities, such as types, cardinality (or multiplicity) of attributes, and keys.

Linguistic resources are used in order to match words based on linguistic relations between them.

Alignment reuse exploit external resources, which record alignments of previously matched ontologies. This technique is motivated by the intuition that many ontologies to be matched are similar to already matched ontologies, especially if they are describing the same application domain.

Upper level and domain specific formal ontologies could be used as external sources of knowledge to help the correspondence identification.

In the same way, structure-level techniques include specific techniques such as:

Graph-based techniques consider the ontologies as labelled graph structures. Usually, the similarity comparison between a pair of nodes from the two ontologies is based on the analysis of their positions within the graphs. The intuition behind this is that, if two nodes from two ontologies are similar, their neighbours must also be somehow similar.

Taxonomy-based techniques are also graph algorithms which consider only the specialisation relation. The intuition behind taxonomic techniques is that is-a links connect terms that are already similar, therefore their neighbours may be also somehow similar.

Repository of structures store ontologies and their fragments together with pair-wise similarity measures. The goal is to identify structures which are sufficiently similar to be worth matching in more detail, or reusing already existing alignments, thus, avoiding the match operation over the dissimilar structures.

Model-based techniques handle the input based on its semantic interpretation. The intuition is that if two entities are the same, then they share the same interpretations.

Data analysis and statistics techniques are those which take advantage of a representative sample of a population. This helps in grouping together items or computing distances between them. Usually, these techniques measure the similarity between sets of annotated entities (e.g., co-occurrence measures).

2.2.1 Matching Evaluation

The growing interest in the ontology matching field has motivated the development of matching systems implementing different matching approaches. With the increasing of algorithms to automatic ontology matching, arose the need of evaluating them. To achieve this, there is an initiative named Ontology Alignment Evaluation Initiative (OAEI) [47]. Since 2004, the OAEI organises campaigns to evaluate matching systems implementing different matching approaches. The goals of this initiative are assessing strengths and weaknesses of matching systems, comparing performance of techniques, and improve evaluation techniques.

In OAEI, the ontologies are described in OWL language and the automatically generated alignments from matchers must be provided in the Alignment API [4] format. The data sets used in the evaluations involve different domains, languages, and sizes addressing the different characteristics that ontologies and alignments could express. The data sets including domains such as anatomy, conference organization, biological/medical classifications, social sciences, among others. Each data set is composed by a set of ontologies in a same domain and usually include reference alignments which are used to automatically evaluate the systems.

The performance of participating matching systems are evaluated through the experimental evaluation of the techniques performances typically adopting quantitative metrics as precision, recall and F-measure which are commonplace measures in information retrieval. In this sense, true positive (TP) could be defined as the correct found correspondences. False positive (FP) are incorrect found correspondences. True negative (TN) are non-existent correspondences correctly classified and false negative (FN) are non-existent correspondences which could be found.

Precision is defined as a set of retrieved data that are relevant (true positive retrieved). Given a reference alignment R, the precision of some alignment A is given by:

$$Precision(A, R) = \frac{TP}{TP + FP}$$

Recall is a set of relevant data that are retrieved from the total of existing correspondences to be found (true positive expected). Given a reference alignment R, the recall of some alignment A is given by:

$$Recall(A, R) = \frac{TP}{TP + FN}$$

F-measure is used in order to aggregate the result of precision and recall. Given a reference alignment *R*, the F-measure of some alignment *A* is given by:

$$F - measure(A, R) = 2 * \frac{Precision * Recall}{Precision + Recall}$$

In this work we will refer later to some of the matching systems that participated in the OAEI campaign, and we consider in our analyses the OAEI Conference data set in particular, the reason for this choice is explained in following chapters.

2.3 Semantic resources and sense disambiguation

An important aspect of our work is the employment of external background knowledge resources in our proposed approach, as discussed later in this thesis. Here we present the basics regarding these resources.

2.3.1 WordNet

WordNet [43, 6] is a general-purpose large lexical database of English frequently adopted as an external resource in automatic ontology matching between domain ontologies [73, 71, 56]. Nouns, verbs, adjectives and adverbs are grouped into sets of cognitive synonyms (synsets), each expressing a distinct concept. A synset denotes a concept or a sense of a group of terms. WordNet also provides textual descriptions of the concepts (gloss) containing definitions and examples. The WordNet 3.0 release has 117,798 nouns, 11,529 verbs, 22,479 adjectives, and 4,481 adverbs. The average noun has 1.23 senses. The main relation among words in WordNet is synonymy and the most frequently encoded relation among synsets is the super-subordinate relation (also called hyperonymy, hyponymy or ISA relation) [6]. WordNet can be accessed on the Web or downloaded and accessed locally. Figure 2.5 shown an example of WordNet senses from the noun Conference with three senses for the noun, each of them has a gloss (a dictionary-style definition), a list of synonyms for the sense, and sometimes also usage examples (shown for the adjective sense).

Figure 2.5 – Example of WordNet senses from Conference

2.3.2 Word Sense Disambiguation

WordNet is a resource widely adopted to represent word senses and relations between senses, for instance, the *is-a* relation between *dog* and *mammal* from WordNet. A sense is a discrete representation of one aspect of the meaning of a word.

Word sense disambiguation (WSD) is the task of determining which sense of a word is being used in a particular context [28]. WSD algorithms take as input a word in context and a fixed inventory word sense disambiguation of potential word senses and outputs the correct word sense in context. The input and the senses depends on the task.

According to Jurafsky and Martin [28], supervised machine learning approaches are often adopted to handle WSD tasks, however, we need to have sufficient data that has been hand-labeled with correct word senses. Moreover, such labeled training data are expensive and limited. Alternatively, the most well-studied dictionary-based algorithm for sense disambiguation is the Lesk algorithm. Lesk algorithm is a family of algorithms that choose the sense whose dictionary gloss or definition shares the most words with the target word's neighborhood.

The Simplified Lesk algorithm computes the overlap and returns the number of words in common between two sets, ignoring function words or other words on a stop list. Figure 2.6 shown an example of simplified Lesk algorithm.

Measuring the minimum editing distance between strings can also be used to identify the sense that better express the meaning of a term. In this way, Levenshtein distance [28] can be applied to determine how similar a string or sequence or characters are. In other

```
function SIMPLIFIED LESK(word, sentence) returns best sense of word

 best-sense ← most frequent sense for word
 max-overlap ← 0
 context ← set of words in sentence
 for each sense in senses of word do
   signature ← set of words in the gloss and examples of sense
   overlap ← COMPUTEOVERLAP(signature, context)
   if overlap > max-overlap then
        max-overlap ← overlap
        best-sense ← sense
 end
 return(best-sense)
```

Figure 2.6 – Simplified Lesk algorithm [28].

words, Levenshtein Distance is defined as the minimum number of operations (insertions, deletions, or substitutions) needed to transform one string into another. Formally, Levenshtein distance version, presented by Jurafsky [28] considers a weight factor in which the insertions and deletions each have a cost of 1 (ins-cost(·) = del-cost(·) = 1), and substitutions have a cost of 2 (except substitution of identical letters have zero cost). Figure 2.7 shown an example of the class of dynamic programming algorithms adopted to calculate the minimum edit distance.

```
function MIN-EDIT-DISTANCE(source, target) returns min-distance

  n ← LENGTH(source)
  m ← LENGTH(target)
  Create a distance matrix distance[n+1,m+1]

  # Initialization: the zeroth row and column is the distance from the empty string
       D[0,0] = 0
       for each row i from 1 to n do
          D[i,0] ← D[i-1,0] + del-cost(source[i])
       for each column j from 1 to m do
          D[0,j] ← D[0,j-1] + ins-cost(target[j])

  # Recurrence relation:
  for each row i from 1 to n do
       for each column j from 1 to m do
         D[i,j] ← MIN( D[i−1,j] + del-cost(source[i]),
                       D[i−1,j−1] + sub-cost(source[i], target[j]),
                       D[i,j−1] + ins-cost(target[j]))
  # Termination
  return D[n,m]
```

Figure 2.7 – Levenstein distance example [28].

2.4 Word Embedding

Word embedding has been largely adopted in several NLP tasks [42]. It is an umbrella name for a set of NLP language modelling and feature learning techniques which represent words as vectors in a semantic space. Models are trained to produce a vector space and reconstruct the linguistic contexts of words. Each unique word in the corpus is assigned a corresponding vector in the space. Word vectors are positioned in the vector space such that words that share common contexts in the corpus are located in close proximity to one another in the space. The similarity between words is calculated using functions as the Cosine similarity and Euclidean distance. Such approach represents an alternative to WordNet similarities, which may fail due to the low WordNet coverage of specific domains.

Different algorithms have been developed to generate embeddings, such as GloVe [53] and Word2vec [42], resulting in a set of available pre-trained models. GloVe is an *unsupervised learning algorithm to obtain vector representations for words*[1]. The training phase uses the Wikipedia 2014 and Gigaword5 corpora. It has 6 billions tokens, 400 thousand vocabulary size and neural network dimension of 200. The GoogleNews model is based on Word2vec and it is trained on part of Google News dataset (about 100 billion words). The model contains 300-dimensional vectors for 3 million words and phrases.

2.5 Conclusions

This chapter has discussed the background related to the topics, resources and techniques adopted in our work. First, we described the fundamentals of ontology, top-level ontologies, and ontology matching. Then, we presented the basis of evaluation matching systems developed by OAEI. After, we detailed the main characteristics of the lexical database WordNet which is adopted in our work, and finally, we summarize Word Sense Disambiguation techniques and Word Embedding models adopted in our experiments. Our approach deals with Lesk and Levenshtein distance to identify WordNet senses that better express the meaning of ontology domain concepts. Alternatively, we adopt Word Embedding models as an alternative way to Word Sense Disambiguation. Chapter 4 details how each technique was implemented in the context of our proposed approach.

[1] https://nlp.stanford.edu/projects/glove/

3. RELATED WORK

Whereas the area of ontology matching [5] has developed fully in the last decades, the problem of matching ontologies involving top-level ontologies has been less studied. However, a growing interest has been shown in the literature in the last few years [33, 57]. One can distinguish different tasks involving these kinds of ontologies, in particular the tasks of matching a pair of top ontologies or a pair involving a source domain ontology with a target top ontology. While the former addresses the problem of matching ontologies with a same level of abstraction, but with their different coverage, conflicting axioms and incompatibilities in restrictions, the latter has to address the challenging task of matching ontologies that may present significant differences in their levels and expressiveness.

Matching top and domain ontologies is a complex task, both automatically and manually. This task requires the deep identification of the semantic of concepts and, in particular, the identification of subsumption relations. In fact, the problem of finding subsumption relations are often neglected by most state-of-the-art matchers. The main problem of matching top and domain ontologies using these currently available matching systems is that, despite the variety of approaches, most of them typically, for an initial estimate of the likelihood that two elements refer to the same real world phenomenon, rely on string-based techniques. Therefore, the found correspondences often represent equivalences with concepts that are equally or similarly written. However, in many cases, this type of correspondence is not the case for ontologies of different levels [58].

Considering this scenario, this chapter reviews the different tasks of ontology matching involving top ontologies: (i) matching of top ontologies; (ii) matching of top ontologies to lexicons; (iii) matching domain ontologies with the help of top ontologies; and (iv) matching top ontologies to domain ontologies.

3.1 Matching top ontologies

As stated in [31], while the purpose of a foundational ontology is to solve interoperability issues among ontologies, the development of different foundational ontologies re-introduces the ontology interoperability problem. Early works addressed this problem [12, 62, 68] on different perspectives. While [12] compared specific treatments of fundamental issues (as significant discrepancies related to universals and particulars, qualities, constitution and spatio-temporality) and how similar notions apply differently in BFO and DOLCE, [62] compared the primitive relations (dependence, quality, and constitution) between these ontologies. In [68], an alignment between BFO and DOLCE was established in order to conciliate their respective realistic and cognitive points of view and to integrate medical data. In [31, 32], alignments between BFO, DOLCE and GFO have been established with automatic matching tools and manually, with substantially fewer alignments found by the matching tools. The alignments in the context of the whole ontology revealed a considerable amount of logical inconsistencies.

In [44], the core characterization of mereotopology of SUMO and DOLCE has been studied, relating their axiomatizations via ontology alignments. This included corrections and additions of axioms to the analyzed theories which eliminate unintended models and characterize missing ones. Finding alignments between DOLCE and SUMO was also addressed in [48], where other ontologies, the SmartDOLCE and SmartSUMO ontologies, have been developed on the basis of DOLCE and SUMO. The alignment of the SUMO taxonomy to DOLCE involved prunning the upper-level of the SUMO taxonomy and the non-trivial task of aligning the remaining concepts to appropriate DOLCE categories.

Aligning higher level ontologies reveals also the problem of matching their different versions. In [61], a method for tracking, explaining and measuring changes between successive versions of BFO1.0, BFO1.1, and BFO2.0 was applied. The aim was to provide a more comprehensive analysis of the changes with respect to the BFOConvert tool[1] which provides an alignment between previous BFO versions.

3.2 Matching top ontologies to lexicons

Several works on equipping lexical resources with top ontologies are developed in order to associate a formal semantics to their lexical layer. As stated in [8] while WordNet has been used in numerous works as an ontology, *"it is only serviceable as an ontology if some of its links are interpreted according to a formal semantics that tell us something about the world and not just about the language"*. In that sense, in the context of the OntoWordNet, they have investigated different ontological problems in WordNet (e.g., confusion between concepts and individuals, constraints violations, heterogeneous levels of generality, etc.). In [8] the WordNet taxonomy with a more rigorous semantics via an alignment between WordNet top-level synsets and DOLCE is provided . This work has been further extended [9] in order to extract association relations from WordNet, and to interpret those associations in terms of a set of conceptual relations in DOLCE. Later, this alignment has been updated [10] with a revision of the manual alignments and a different version of DOLCE and WordNet. While these works focused mostly on WordNet noun synsets, [64] extended the previous alignments by aligning verbs according to their links to nouns denoting perdurants, transferring to the verb the DOLCE class assigned to the noun that best represents that verb's occurrence. They argue that many NLP applications need to deal with events, actions, states, and other temporal entities that are usually represented by verbs.

[1] http://ontobull.hegroup.org/bfoconvert (last viewed on March 25th, 2019)

In [60], a semi-automatic method for aligning WordNet3.0 and BFO2.0 is described. It adopts previous alignments between WordNet and the KYOTO ontology, whose top layer is based on DOLCE. The method involves manually creating a set of alignments between the ontologies and implementing a set of matching rules. The manual creation of the alignments explores diverse existing ones: a) KYOTO and BFO (on the basis of previous alignments between DOLCE to BFO1.0 and BFO1.1 [62, 68, 31], ignoring the axiomatization incompatibilities); b) BFO1.0 and BFO1.1 to BFO2.0 (on the basis of the alignments in [61]); and c) WordNet labels and BFO2.0.

In [55], the authors report the matching and integration of several background resources and ontologies of varying complexity to the Cyc knowledge base. These resources and ontologies included large pharmaceutical and medical thesauri and large portions of WordNet. For this task, ontologists have been trained with domain experts and interactive clarification dialog-based tools were developed to enable experts to directly match/integrate their ontologies.

In [46], SUMO has been aligned to WordNet 1.6[2]. The applied methodology identified relations of synonymy, hypernymy, and instantiation of SUMO concepts with WordNet synsets, where higher level notions in WordNet were mapped to an equivalent in SUMO. Finally, in [34], WordNet has been extended by applying the notion of *semantic types* in order to establish matching rules between the noun synsets of Wordnet and the top-level constructs of the UFO ontology.

3.3 Matching domain ontologies via top ontologies

As top ontologies provide a reference for rigorous comparisons of different ontological approaches, they also serve as a framework for analysing, harmonizing, matching and integrating existing domain ontologies [48]. In domain ontology matching, in particular, they may act as semantic bridges. Despite the potential gain of exploiting top ontologies in domain ontology matching, few works have addressed this alternative, possibility due to the still low coverage of top ontologies in domain ontologies. This gain has been quantitatively measured in [39], where a set of algorithms using such semantic bridges are applied. The circumstances of cases where top ontologies improve domain ontology matching, with respect to approaches ignoring them, were then studied. The experiments were run with SUMO-OWL (a restricted version of SUMO), OpenCyc and DOLCE and demonstrate that overall the alignment via upper ontologies impacts in F-measure positively. Additionally, in [49] a set of alignment patterns based on the OntoUML (a conceptual modelling language

[2]With other versions of WordNet been made available later https://github.com/ontologyportal/sumo/tree/master/WordNetMappings

based on UFO) are applied to a set of alignments generated by matching systems. An analysis of the impact of the use of the patterns to avoid common errors was presented.

Very few concrete matching approaches however exploit top ontologies. An example is the semi-automatic LOM matcher [37], which applies four methods (1) whole term matching; (2) word constituent matching; (3) synset matching; and (4) type matching. Type matching explores the ontological category of each word constituent for matching using the alignments from WordNet synsets to SUMO. LOM takes the source terms that are unmatched with the three first methods, collects the set of SUMO terms that their synsets map to, and then compares the SUMO term sets to their counterpart for each term in the target ontology.

From a manually established alignment between biomedical ontologies and BFO, in [63], a matching approach relies on filtering out correspondences at domain level that relate two different kinds of ontology entities. The matching approach is based on a set of similarity measures and the use of top ontology as a parameter for better understanding the conceptual nature of terms within the similarity calculation step. Besides the improvement in the results obtained, the introduction of top ontologies in the alignment process increased the influence of semantic factors in this task, further expanding the universe of information to be explored during the alignment.

3.4 Matching domain ontologies to top ontologies

Methodologies for constructing ontologies should not neglect the use of top ontologies and should better address it in a *top-down* approach [30]. In the absence of the systematic use of foundational ontologies within domain ontology development, a *bottom-up* approach has to be applied instead. Few automatic solutions have been however proposed in the literature. In [1], DOLCE has been used to integrate two geoscience knowledge representations, the GeoSciML schema and the SWEET ontology, in order to facilitate cross-domain data integration. The aim was to produce a unified ontology in which the GeoSciML and SWEET representations are aligned to DOLCE and to each other. In that perspective, DOLCE works as a semantic bridge and this approach fits as well the category of domain matching with foundational ontologies. The alignments have been manually established and representation incompatibility issues have been discussed so far. DOLCE has been manually aligned to the domain ontology describing services (OWL-S) in [41], in order to address its conceptual ambiguity, poor axiomatization, loose design and narrow scope. They have also developed a core ontology of services to serve as middle level between the foundational and OWL-S, and can be reused to align other Web Service description languages.

In [3], several schemata of FactForge, which enables SPARQL query over the LOD cloud, have been aligned to the foundational ontology PROTON in order to provide a unified way to access to the data. The alignments were created by knowledge engineers through a systematic process. Equivalence (e.g., Geonames:Country equivalentClassOf PROTON:Country) and subclass relationships (e.g., DBpedia:OlympicResult subClassOf Proton:Situation) between DBPedia, Geonames and Freebase concepts and PROTON classes have been established. As stated in §3.3, manually alignments have also been established between biomedical ontologies and BFO, in [63].

While these proposals mainly generate manual alignments between foundational and domain ontologies, one of the few automatic approaches is BLOOMS+ [27]. It has been used to automatically align PROTON to LOD datasets using as gold standard the alignments provided in [3]. BLOOMS+ first uses Wikipedia to construct a set of category hierarchy trees for each class in the source and target ontologies. It then determines which classes to align using 1) similarity between classes based on their category hierarchy trees; and 2) contextual similarity between these classes to support (or reject) an alignment. BLOOMS+ significantly outperformed existing matchers in the task.

Table 3.1 – Summary of the approaches on chronological order

	Foundational/lexicon/domain	*Approach*	*Available alignment*
Matching top ontologies			
[12]	BFO1.0, DOLCE	Manual comparison	-
[48]	SUMO, DOLCE	Manual alignment	-
[62]	BFO1.0, DOLCE	Manual comparison	-
[68]	BFO1.0, DOLCE	Manual alignment	Set of triples
[31]	BFO1.1, DOLCE-Lite, GFO	Manual, matching tools	List at Romulus[1]
[61]	BFO1.0,1.1,2.0	Semi-autom. (change-tracking)	-
[44]	SUMO, DOLCE-CORE	Manual alignment	FOL alignments
Matching top ontologies to lexical resources			
[8, 9]	DOLCE-LitePlus,DOLCE-UltraLite/WordNet1.6	Semi-autom. (NLP, disamb., A-links)	OWL version[2]
[55]	Cyc/WordNet1.6	Semi-autom. (interactive tool, rules)	-
[46]	SUMO/WordNet1.6/3.0	Manual	Textual format
[10]	DOLCEPlusDnS Ultra Lite/WordNet3.0	Semi-autom. (transitive closure)	RDF dataset
[60]	BFO2.0/WordNet3.0	Semi-autom. (matching rules)	-
[64]	DOLCE-LitePlus/WordNet3.0 (verbs)	Semi-autom. (annotation tool, links)	-
[34]	UF0/WordNet3.0	Automatic (SemanticMapper)	-
Matching domain ontologies via top ontologies			
[37]	SUMO, Cyc/SENSU	Semi-autom. (LOM matcher)	-
[39]	SUMO-OWL, OpenCyc, DOLCE/17 ont.(agent, bibtex, etc)	Automatic (structural matching)	-
[63]	BFO/GO, INOH Event	Automatic (FOAM+OBOAEA)	-
[49]	UFO/Conference	Manual pattern analysis	-
Matching domain ontologies and top ontologies			
[41]	DOLCE/OWL-S	Manual	-
[1]	DOLCE-LitePlus/GeoSciML2.0, SWEET1.1	Manual	UML-syntax
[3]	PROTON/DBPedia, Freebase, Geonames	Manual	-
[27]	PROTON/DBPedia, Freebase, Geonames	Automatic (BLOOMS+)	-
[63]	BFO/GO, INOH Event	Manual	
Our work	DOLCE Ultralite, SUM0/Conference, SSN, CORA	Automatic (WordNet + Embeedings)	OAEI API

3.5 Discussion

Table 3.4 summarises the matching approaches involving foundational ontologies described in this paper. Most approaches still rely on manually or semi-automatically established alignments. This task is far from being trivial, even when done manually. This has been recently corroborated in [66], where manually classifying domain entities under top ontology classes is reported to be very difficult to do correctly. On top of that, manual ontology matching is also an expensive task that may introduce a bias as it represents a point of view expressing the interpretation of the concepts influenced by the background of the expert. As knowledge on top ontologies is highly specialized, it is important that such evaluation considers an overview of different experts in this area. Moreover, while manual alignment on a small set of concepts is feasible, bigger data sets would require more efforts. The findings in [66] also point out the need for improving the methodological process of manual integration of domain and foundational ontologies, in accord with what has been stated in [30].

Systematically equipping domain ontologies with foundational ones would also promote their use as semantic bridges [39, 63] in the task of matching domain ontologies. Despite the variety of approaches focusing on domain ontologies, few works exploiting top ontologies as bridges have been proposed in the literature.

While more automation is an obvious requirement in the field, the poor performance of solutions addressing automatically matching different top ontologies or with domain ontologies have demonstrated the difficulty of the task, as reported in experiments evaluating currently matching tools [31, 58]. Current tools fail on correctly capturing the semantics behind the ontological concepts, what requires deeper contextualization of the concepts on the basis of their hierarchy and axioms. In that sense, further context and documentation is required, in particular for domain ontologies, to help identifying the right semantics behind the context (e.g, the ontologies from the largely used OAEI Conference dataset have a very poor lexical layer). Besides that, the task requires the identification of other relations than equivalences, such as subsumption and meronym. The latter is largely neglected by current matchers. In fact, when having different levels of abstraction it might be the case that the matching process is capable of identify subsumption correspondences rather than equivalence, since the top ontologies have concepts at a higher level. Furthermore, while diverse matching approaches rely on external background knowledge (BabelNet, WordNet, UMLS, etc.), the coverage of top ontologies in these resources is still low. More recently, the resource Framester[3], exposed as a knowledge graph, addresses this aspect as a hub between several resources such as VerbNet, BabelNet, DBpedia, Yago and DOLCE-Zero.

While the automatic approaches have been mostly manually evaluated, with few exceptions [3, 57], systematically evaluations of matching systems have been so far dedicated to domain ontologies. Despite the variety of tasks in the OAEI campaigns[4], the evaluation of matching involving top has not been addressed. Producing comprehensive evaluation data sets on which matching solutions can be evaluated would foster the development of approaches involving top ontologies and support a next generation of semantic matching approaches. With that respect, few of the established alignments generated by the ap-

[3]https://lipn.univ-paris13.fr/framester/
[4]http://oaei.ontologymatching.org/2018/

proaches have being publicly made available, as seen in Table 1. Furthermore, very few of them adopted a format that can be processed by automatic tools.

Another aspect refers to the evolution or the consistency of alignments with respect to the evolution or the different variants of the ontologies. For example, DOLCE and its different variants have been used in diverse proposals, as many efforts have been dedicated to the development of this ontology. DOLCE has been exposed with reduced axiomatization and extensions with generic or domain plugins, such as for DOLCE-Lite [9], DOLCE-Lite-Plus[5] or still DOLCE+DnS Ultralite[6]. Besides their substantial differences in the hierarchical organization and expressiveness, these versions are mostly compatible, what is not the case for other ontologies. For instance, BFO 2.0 represents major updates to BFO not strictly backwards compatible with BFO 1.1 and a manual alignment was required to express their incompatibilities. Evolving alignments to cope with the different versions of the ontologies is still an open challenge.

Last, but not least, very few foundational ontologies are equipped with lexical layers in other natural languages than English (e.g., BFO has been enriched with a lexical annotation in Portuguese). However, with the increasing amount of multilingual data on the Web and the consequent development of ontologies in different natural languages, foundational ontologies should also be equipped with richer multilingual annotations in order to facilitate the multilingual and cross-lingual ontology matching tasks.

Considering that the state of the art lacks fully automated alignments for top and domain ontologies and the high complexity of the task, in this thesis we propose an automatic approach for matching domain and top ontologies that exploits existing alignments between semantic resources (WordNet) and top ontologies. For that purpose our work consider two top ontologies and three domain ontologies as discussed in detail in the other chapters. Our results are provided following the OEAI API standards.

3.6 Conclusions

This chapter has presented an overview of the different tasks involved in matching foundational ontologies. We have discussed the weaknesses of existing proposals and highlight the challenges to be addressed in the field. Being know as a very hard task and with very few automatic approaches, we propose a first evaluation of a matching system devoted to generated domain and top ontology alignments that is based on previous given alignments of top ontologies with semantic resources. As such evaluation has not been made before, to the best of our knowledge, we considered that worth investigating. This approach is introduced in the next chapter.

[5]http://www.loa.istc.cnr.it/old/ontologies/DLP_397.owl
[6]http://www.ontologydesignpatterns.org/ont/dul/DUL.owl

4. PROPOSED APPROACH

This chapter presents our proposed approach to align domain and top-level ontologies. Despite the variety of approaches to match ontologies, they have been developed mainly focusing in domain ontologies. In a previous work [58], we could observe that current matching tools can not deal with the task of matching domain and top-level ontologies. In order to provide an approach to the task of aligning domain and top-level ontologies, we propose an approach that is based on existing external resources as presented below.

The matching process in our approach is divided in three main steps. The first step corresponds to a pre-processing phase. The second step disambiguates the domain concept with respect to the entry in the external resource (i.e. finding the appropriate WordNet synset representing the domain concept, when exploring existing alignments between WordNet and top-level ontologies); the third step finds the alignments between the domain and top-level concepts. Figure 4.1 presents a general view of the proposed approach. In the next sections, we detail each step.

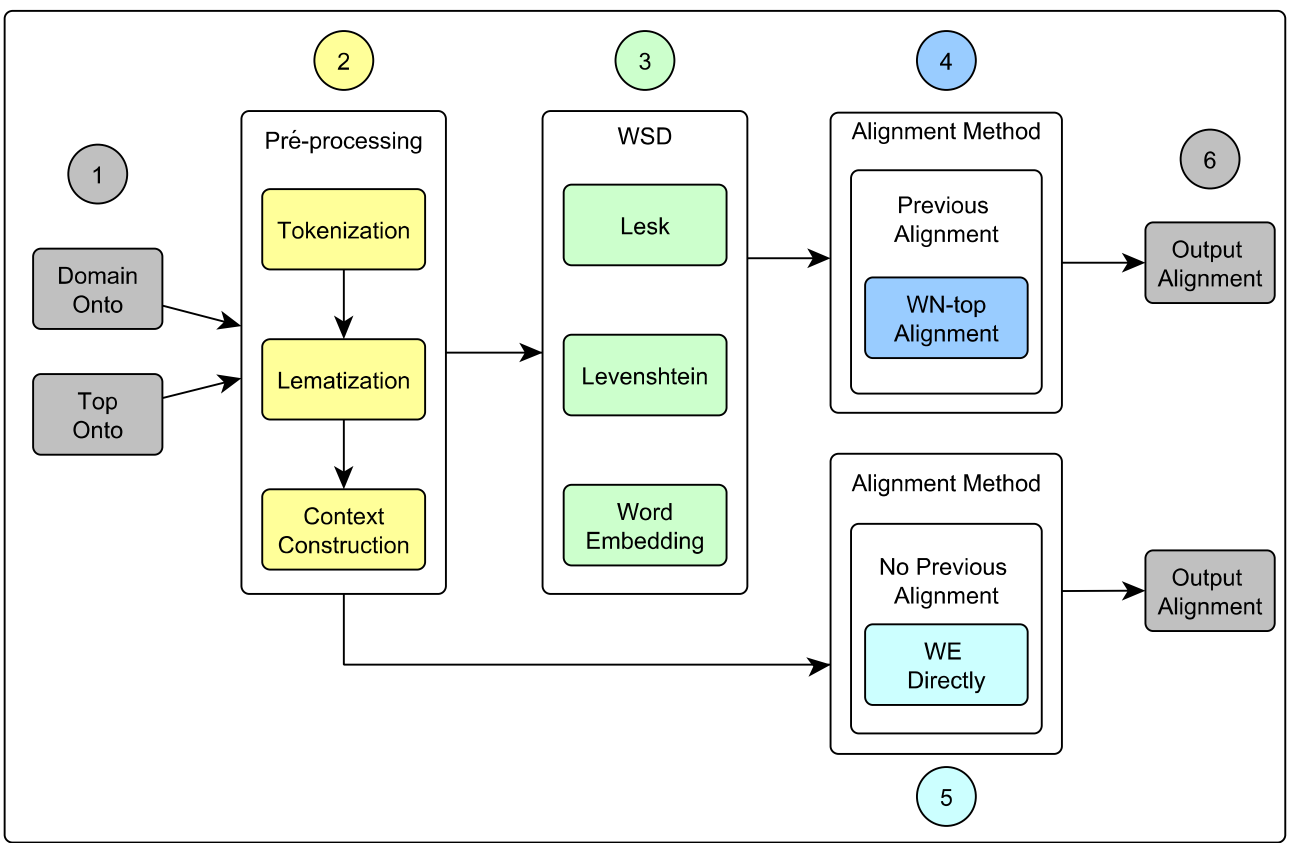

Figure 4.1 – General view of the proposed approach.

4.1 Pre-processing

After selecting the domain and top-level ontologies which will be aligned, our approach starts with a pre-processing phase. All information available about an ontology entity, including entity naming (ID), annotation (usually labels and comments) and information on the neighbours (super and sub-concepts) are tokenized and lemmatized, then, stop words and special characters are removed, creating a *context* of each ontology concept.

Given *Sup*(e) and *Sub*(e), the sets of terms denoting the super-concepts and sub-concepts of the entity e, and *Ann*(e) the set of terms from its annotations, a naive strategy for building a *context* (*context*) considers these sets as a *bag of words*. A formal representation of the context is available below:

$$context(e) = \{e, w | w \in Sup(e) \cup w \in Sub(e) \cup w \in Ann(e)\}$$

Alternatively, we also adopt all ontology concepts as a *bag of words* (the alternative configuration will be more evident in Chapter (6). The notion of *context* is adopted in order to select the closer sense of the domain concept in the semantic database WordNet [6]. In this way, we adopt three Word Sense disambiguation strategies: *Lesk measure*, *Levenshtein distance*, and *Word embedding*. In the next section these strategies are detailed.

4.2 Word Sense Disambiguation

Disambiguation is the process of identifying a meaning of a word in a given context. As mentioned above, our approach considers the *context* composed by a bag of words from the ontology information to select a sense that represents the meaning of the given concept in a semantic database (WordNet). In this way, we retrieve all available meaning from the domain concept in the lexical database and, if the term is monosemic (there is only one sense), it is assigned as the sense for the concept. However, if the term is polysemic (there are more than one sense), we start the disambiguation process.

In order to perform the disambiguation, each retrieved WordNet sense, is also transformed in a bag of words composed by its terms and gloss. Similarly to the *context* of ontology concepts, here we create a *context* of each sense, considering all information of them, after tokenize, lemmatize, and remove stop words and special characters. A formal representation of the context of WordNet senses is available below:

$$context(s) = \{t_1, t_n \cup Gloss(t)\}$$

Where, s corresponds to the WordNet senses and t represent the terms related to the sense.

In some cases, the domain concept is composed by more than one word, hence, first we look on the database if the composed term exists. If there is no correspondence, we adopt the last part of the concept term (assuming that this is his head noun).

Our approach deals with Lesk and Levenshtein distance to identify WordNet senses that better express the meaning of ontology domain concepts. Alternatively, we adopt Word Embedding models as a complementary way to Word Sense Disambiguation. Next, we discuss each of the Word Sense disambiguation strategy considered in our approach.

4.2.1 Lesk Measure (Exact match)

The Lesk measure for word sense disambiguation [35] relies on the calculation of the word overlap between the sense definitions of two or more target words. Given a word w, it identifies the sense of w whose textual definition has the highest overlap with the words in the context of w :

$$score_{Lesk}(S) = |context_{Lesk}(w) \cap gloss(S)|$$

where $context_{Lesk}(w)$ is the bag of all content words in a context window around the target word w. Here, we overlap the $context(e)$ with the context of each WordNet synset $context(synset)$:

$$context(synset) = \{w|w \in Terms(synset) \cap w \in Gloss(synset)\}$$

where $Terms(synset)$ the set of terms in a $synset$ and $Gloss(synset)$ the corresponding set of terms from the gloss (i.e, textual description containing definitions and examples) associated to the synset. We hence retrieve the highest overlap between $context(e)$ and $context(synset)$

$$score'_{Lesk}(e) = |context(e) \cap context(synset)|$$

4.2.2 Levenshtein Distance (String distance)

The Levenshtein distance for word sense disambiguation [36] relies on the calculation of the minimum number of single-character edits (i.e. insertions, deletions or substitutions) required to change one word into the other. Hence, we compute the cost of edition to transform one string into another.

Given two words w_s and w_t, it computes the weight of edition in which the shortest path of edition is needed from one string to another. Here, we compare the *context*(*e*) with the context of each WordNet synset *context*(*synset*) retrieving the string distance of each term. The synset with the lowest edition cost is selected.

$$score'_{Levenshtein}(e) = |cost_{context(e)} \cap cost_{context(synset)}|$$

4.2.3 Word Embedding

In this step, we adopt previous trained Word Embedding models as a way to retrieve a similarity of the context of the domain entities *context*(*e*) with the context of the WordNet synsets *context*(*synset*) as a strategy of Word Sense Disambiguation. In our approach we consider two available Word Embedding models (GloVe: Global Vectors for Word Representation [53] and Word2vec [42]). The models present a similarity weight of strings which were considered to identify the sense that better represents the meaning of the domain ontology concept.

GloVe is an unsupervised learning algorithm for obtaining vector representations for words. Training is performed on aggregated global word-word co-occurrence statistics from a corpus, and the resulting representations showcase interesting linear substructures of the word vector space. In our work, we adopt the pre-trained word vectors model based on Wikipedia 2014 + Gigaword 5 (6B tokens, 400K vocab, uncased, 50d, 100d, 200d, and 300d vectors).

In the same way, Word2vec tool takes a text corpus as input and produces the word vectors as output. It first constructs a vocabulary from the training text data and then learns vector representation of words. This method adopts the cosine distance between two words generated by the word embedding model to identify the similarity between them. There are two main learning algorithms in Word2vec: continuous bag-of-words and continuous skip-gram. Both algorithms learn the representation of a word that is useful for prediction of other words in the sentence. In our work, we adopt the pre-trained word vectors model trained on part of Google News dataset (about 100 billion words). The model contains 300-dimensional vectors for 3 million words and phrases.

As mentioned, we adopt the trained models to retrieve the similarity (*sim*) between strings from *context*(*e*) and *context*(*synset*), then we calculate the average similarity. First, we calculate this average to all elements of the context, individually, then, we calculate the average of the context, considering the context length. The synset with the higher average is selected.

$$score'_{WE}(e) = |sim_{context(e)} \cap sim_{context(synset)}|$$

4.3 Alignment Method - via previous WordNet alignments

In the steps described above, we pre-process the information available on the ontologies and adopt one of the word sense disambiguation strategy to select the WordNet synset considered closer to the meaning of each domain concept. In this step we perform the identification of the top-level concept. This task relies on the representation of the given existing alignments of WordNet and the top-level ontologies DOLCE and SUMO.

4.3.1 DOLCE correspondence identification

This step uses existing alignments between DUL and WordNet 1.6 (OntoWordNet). In OntoWordNet, the WordNet hyponymy relation is aligned to the subsumption relation and the synset notion could be aligned to the notion of concept. The named concepts were normalized to obtain one distinct name for each synset. Hence, if a synset had a unique noun phrase, it is used as a concept name (*e.g.* Document__Written_Document__Papers). If the noun phrase was polysemous, the concept was numbered (*e.g.* Writting_1, Writting_2). Figure 4.2 presents a fragment of WordNet synsets (as concepts) linked to DUL concepts. The first-level concepts (in lower case) correspond to a DUL concept. The upper case concepts represent WordNet synsets. Each concept in OntoWordNet is associated to an annotation containing the corresponding gloss of the synset in WordNet.

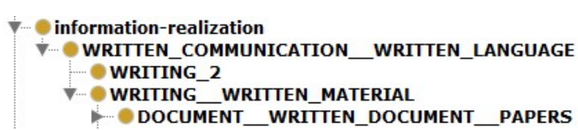

Figure 4.2 – Example of WordNet synsets linked to DOLCE.

For each concept of the domain ontology, we use the selected synset (Section 4.2) to identify the corresponding concept in OntoWordNet. To select the concept in OntoWordNet we compare the WordNet synset with each concept c in OntoWordNet (recall that concepts are represented by the concatenation of words). A bag of words for the OntoWordNet concept is created from the concatenated words and gloss, i.e., $context(c)$. Then, we overlap the synset and c.

$$score'_{WSD}(c) = |context(c) \cap context(synset)|$$

After finding the OntoWordNet concept c corresponding to the synset, the higher level concept h^c of c is retrieved,. h^c corresponds to the DOLCE concept.

4.3.2 SUMO correspondence identification

Similarly to the correspondence identification in DOLCE, this step uses existing alignments between SUMO and WordNet 3.0, in order to identify the domain and top concepts correspondences. For each identified correspondence, the synset of WordNet is augmented with three information : *(i)* a prefix (&%) that indicates that the term is taken from SUMO; *(ii)* the SUMO concept; and *(iii)* a suffix indicating the kind of relation. The suffix '=' indicates that the correspondence relation is synonymy. '+' indicates that the concept is a hypernym of the associated synset. The instantiation relation is indicated by the suffix '@'. An example of the structure of a correspondence representing a synonymy relation can be seen below. In the example, *"06470073 10 n 03 document 0 written_document 0 papers"* corresponds to the synset. The gloss is defined as *"writing that provides information (especially information of an official nature)"*, the prefix *"&%"* indicates that the term is taken from SUMO. *"FactualText"* corresponds to the SUMO concept and the signal *"+"* is the suffix indicating the hyponymy relation.

06470073 10 n 03 document 0 written_document 0 papers | writing that provides information (especially information of an official nature) &% FactualText+

As SUMO-WordNet alignment is a file containing the synset ID, terms, gloss, and the alignment to top concept we search for the domain selected synset in this file and, if the synset is found, we match the domain concept with the top-level concept related to the synset.

As described above, our approach depends on the availability of alignments between the background knowledge resource (here, WordNet) and the top-level ontologies. Hence, we are able to exploit other top-level ontologies in case such alignments exist. This leads also to the question on the maintenance of these alignments with the evolution of the ontologies and the given resource, which is out of the scope of this work. Hence, we alternatively developed some approaches that are independent of the previous WordNet-Top ontologies alignments, as presented below.

4.4 Alignment Method - without previous alignments

In this step, we adopt the word sense disambiguation techniques described in Section 4.2, however, we are not limited to the previous alignments of WordNet-Top level ontologies. Here we investigate the use of the similarity from Word Embedding models to select the correspondences of domain and top ontologies directly.

In this strategy, we consider the context of the domain concept $context(e)$ and, instead of the WordNet synsets, we create a context for top-level concepts $context(top)$. Then, we retrieve the similarity between $context(e)$ and $context(top)$, then, we calculate the average of the $context(top)$ considering the context length. Each top-level concept receives a similarity value, from the values retrieved of the word embedding model. After, we match the domain concept with the top concept with the higher similarity. The word embedding models adopted are the same mentioned in Section 4.2.3

In this way, we do not depend on the availability of the WordNet-Top level ontologies alignments, moreover, it is possible to adopt other trained word embedding models (or training a model for a specific domain) and other top-level ontologies. As we are interested in compare the performance of the previous WordNet top-level alignments, our experiments were developed considering the same trained models and set of ontologies.[1]

4.5 Technical Aspects

In this section, we describe some technical aspects and resources adopted to implement our approach and evaluation. The prototype was implemented in Java programming language. The choice was motivated by the fact that there are ontologies alignment resources available in this language. We work with ontologies in OWL format, then OWL API 5.1.0 [26] was adopted to read and manipulate the ontologies. Regarding the alignment, we use Alignment API 4.0 [4] for accessing the common format. It can be adopted for storing, finding, and sharing alignments, generate tests and comparing alignments (precision, recall), among others. The developed prototype is available on: https://github.com/danielasch/Ontology-Matcher-WN.

[1] Computationally, we observe that the cost of adopt word embedding models instead of WordNet approach is higher. We are not computing the complexity and time to generate the alignments because we are interested mainly in the analysis of the correspondences, however, it is interesting to analyse in this aspect.

4.6 Conclusions

This chapter has discussed our proposed approach to align domain and top-level ontologies. We described the phases adopted in order to identify correspondences between domain and top-level ontologies. First, we detailed the adopted pre-processing steps, next, we discussed the word sense disambiguation techniques which were adopted to select the sense that better express the meaning of the domain concept in WordNet. The Word Sense disambiguation techniques were implemented in parallel, allowing to select which technique will be used in each execution. Complementary, we implemented an alignment method independent of the previous WordNet alignments to top-level ontologies. In this method, we consider similarity through word embedding models as a way to identify correspondences of domain and top-level concepts.

5. REFERENCE ALIGNMENTS

This chapter presents the reference alignments used in our study. The ontologies from the OAEI Conference track have been aligned to the SUMO and DOLCE top-level ontologies. The Conference dataset is one of the most used dataset in ontology matching evaluation and has been extended in several versions. However, it lacks in alignments to top-level ontologies. As a complete manual alignment between SUMO and WordNet is available, we use such alignments as bridges to facilitate the matching task. In the same way, there are alignments of DOLCE (DUL) and WordNet that were also adopted. In Section 5.1, we detail SUMO alignment, in 5.2, we describe the alignments with DOLCE.

Moreover, we describe previously available domain-top alignments (5.3). These are based on domain ontologies developed in a top-down approach considering top-level ontologies as a starting point to specialize the domain concepts. These aligned ontologies were also adopted in our experiments.

5.1 Aligning Conference Ontologies with SUMO

This section presents the alignment of the Conference ontologies [74] to the SUMO [45, 51] top-level ontology. The Conference dataset has been used in the Ontology Alignment Evaluation Initiative (OAEI)[1], which have been carried out over the last fifteen years. The OAEI, however, still lacks matching tasks involving foundational ontologies. The choice for aligning this OAEI dataset to a foundational ontology is motivated by the fact that it offers expressive and real-world ontologies and has became one of the most used in matching evaluations [74]. This dataset has also been extended with different proposals [2, 40], and recently it is also covering complex alignments [70].

We present the process of establishing a consensual alignment between the Conference ontologies and SUMO. Matching foundational and domain ontologies is far from being a trivial task and most approaches still rely on manual or semi-automatic strategies. This has been corroborated in [66], where manually classifying domain entities under foundational ontology concepts is reported to be very difficult to do correctly. As knowledge on foundational ontologies is highly specialized, it is important that such alignments consider the participation of different experts in the area. The findings in [66] also point out the need for improving the methodological process of manual integration of domain and foundational ontologies, in accord with what has been stated in [30]. As a complete manual alignment between SUMO and WordNet has been previously provided [46] and continually updated

[1] http://oaei.ontologymatching.org/

since the original effort, we argue here that using these alignments as bridges to matching domain ontologies to SUMO can facilitate the matching task.

We have chosen SUMO for several reasons. It is the only formal ontology that has a complete set of manually-performed correspondences to all 117,000 word senses in WordNet. It is also one of the few ontologies that has a detailed formalization in an expressive logical language. Most ontologies are still simple taxonomies and frame systems, and so assessing the meaning of their terms requires human intuition based on term names and relationships. SUMO includes a computational toolset [52] that allows users to test the logical consistency of its definitions, which provides a guarantee of quality and correctness than just testing type constraints. Lastly, SUMO is large and comprehensive at roughly 20,000 terms and 80,000 hand-written logical axioms, exceeding the size of other open source foundational ontologies by several orders of magnitude.

We describe the design choices and methodology followed for constructing this bridged alignment. We discuss the main issues that the experts were faced with during the process.

5.1.1 SUMO Alignment Methodology

This section describes the overall methodology we have followed to create the consensual alignment between SUMO and the domain ontologies. As stated above, this process relies on the existing alignment between SUMO and WordNet. We consider these previous alignments as correct so the aim here is to identify the right WordNet synset to the domain concepts. Also, we have reduced the problem to the first-level concepts of the hierarchies from the domain ontologies. This has resulted in 70 first-level domain concepts. For each first level concept of the domain ontology, a foundational specific concept is associated. The cost of doing manual alignment with first level concepts is smaller, as it is reduced to the number of concepts at the first level.

Four evaluators have been involved in the task of aligning the 70 top-level domain concepts to the WordNet synsets. The evaluators are researchers, therefore all have common-sense knowledge about conferences (the domain ontology), they have background in Computer Science and are well acquainted with ontology matching. One of the evaluators is the creator of the SUMO ontology.

The overall methodology is articulated in the following two steps: i) Individually generating the alignments between domain concepts and WordNet synsets; ii) Collaboratively validating the set of found correspondences. Next we detail each step.

5.1.2 Individual Generation of Correspondences

In this first step, each domain concept and the corresponding WordNet synsets, resulting from searching in WordNet for the term associated to the domain concept, were listed to the evaluators. In the absence of entries in WordNet for the terms, a head modifier strategy has been applied (i.e., WrittenPaper is a Paper). Only one concept had not corresponding entry in WordNet (*sigkdd#Sponzor*[2]). In order to help the evaluator to understand the context of the domain concept, their sub-concepts were also presented. As the domain ontologies are equipped with very few comments or labels, we have completed the description of the concept using the definitions from the Cambridge Dictionary[3]. However, we are aware that the found definitions may not reflect the exact semantic of the concept. Each evaluator then was asked to select the right WordNet synset for each domain concept. The evaluators were instructed to select one option for each domain concept, however, in some cases more than one sense was selected. This happens because the domain concept was not clear enough, or the senses available in WordNet were very general. Evaluators were also invited to comment their decisions. Table 5.1 shows a fragment of the spreadsheet for the domain concept *cmt#Conference*.

Table 5.1 – Example of spreadsheet adopted by the evaluators.

Domain concept	WordNet synsets	SUMO concept	Comments
cmt#Conference	**S1**: conference group discussion a discussion among participants who have an agreed (serious) topic **S2**: league conference an association of sports teams that organizes matches for its members **S3**: conference a prearranged meeting for consultation or exchange of information or discussion (especially one with a formal agenda)	S1: Communication+ S2: SportsLeague+ S3: FormalMeeting=	An academic conference would be a FormalMeeting

5.1.3 Validating the Correspondences

After the individual annotation of each domain concept with the WordNet synset, the annotators were able to see the annotations of each other and identify the conflicts. Based on the views on the other annotators (and their comments), each one was able to change their initial annotations. For those conflicts where the comments were not enough for understanding the annotation, an online discussion took place. From the 70 annotated domain concepts, 8 of them have been annotated with different WordNet synsets. Table 5.2

[2]http://oaei.ontologymatching.org/2018/conference/data/sigkdd.owl
[3]http://dictionary.cambridge.org/us/

lists some examples, which are discussed in the next section. All generated alignments are available in the Alignment API format[4].

Table 5.2 – Conflicts solved after discussion between annotators. Boldface concepts represent the conflicts and final results are indicated with a underline.

Domain concept	WordNet synsets	SUMO concept
cmt#Decision	S1: decision determination conclusion S2: decision determination conclusion S3: decisiveness decision S4: decision S5: decision	1: **Learning+** 2: **Deciding+** 3: TraitAttribute+ 4: ConstantQuantity+ 5: ConstantQuantity+
cmt#Preference	S1: predilection preference orientation S2: preference S3: preference predilection taste S4: preference druthers	1: **IntentionalRelation+** 2: SubjectiveAssessmentAttribute+ 3: PsychologicalAttribute+ 4: **PsychologicalAttribute+**
edas#PersonalHistory	S1: history S2: history S3: history account chronicle story S4: history S5: history	1: **PastFn=** 2: History= 3: **HistoricalAccount+** 4: PastFn= 5: Proposition+
sigkdd#Award	S1: award awarding S2: award accolade honor honour laurels S3: prize award	1: UnilateralGiving+ 2: **ContentBearingObject+** 3: **UnilateralGiving+**
conf#ReviewPreference	S1: predilection preference orientation S2: preference S3: preference penchant predilection taste S4: preference druthers	1: **IntentionalRelation+** 2: SubjectiveAssessmentAttribute+ 3: PsychologicalAttribute+ 4: **PsychologicalAttribute+**

5.1.4 Discussion

During the process of alignment construction, several difficulties arose for interpreting the real meaning that the concept represents in the domain ontology. For instance, the concepts `Bid` and `Preference` (Table 5.2) in `cmt` ontology had no description clarifying its use, and no sub or super concepts which could be used to clarify their meaning. In these cases, the evaluators discussed and considered the proper meaning according to their own interpretation of the domain, however, such cases may interfere with the quality of the resulting reference alignment because there is no objective standard for what the meaning, and therefore the correct mapping must be. We have only the consensus guess about intended meaning among human evaluators. In addition, some concepts represented in the ontology present other kind of problems such as doubts regarding ontology elements' adequacy, for example, the concept `ReviewRating` in `edas` ontology, which according to the discussion raised by the evaluators, a rating could be a relationship between a thing, an agent and a rating value. In the same way, the concept `Deadline` in `sigkdd` ontology could be a relationship between the conference and a date. They are however defined in those ontologies as concepts, rather than relationships.

[4] https://github.com/danielasch/ReferenceAlignment

In other cases, sub-concepts are different from first-level concepts and therefore they represent different information, as the concept `Event` in `ConfOf` ontology. Some of their sub-concepts `SocialEvent/Banquet`, `WorkingEvent/Conference`, `WorkingEvent/Workshop` are in line with the main concept, however other such as `AdministrativeEvent/CameraReady Event` seems out of the context. In fact, it should not be a Process at all but a deadline for doing something (submitting a version of a paper, for instance).

In contrast, one can examine a SUMO definition of a term such as `FormalMeeting`[5] and see that it is necessarily a `Meeting` that is not a `SocialParty`, that it must be temporally preceded by a `Planning` that has the `result` of creating the meeting, as well as constraints that other events like a `Resolution` to be considered such, may only occur at a `FormalMeeting`. Something like a modern dictionary, but with the definitions expressed in logic, rather than human language, so that a machine can perform computation (and consistency checking) with those definitions. The cases described above consist of ontological representation problems commonly present in lightweight ontologies, and hinder the reuse and reliability of the represented knowledge. In addition, they highlight the importance of advancing in research that uses top-level ontologies to give more formalization to domain ontologies.

The challenges in aligning the OAEI ontologies should highlight two elements that are lacking in the majority of most of current ontology practice. The first element is the degree of reuse. Ontologies that are created from scratch suffer from the fact that their terms have only a small number of relationships to other terms. The point of having an ontology is to have a shared meaning among its users. When domain ontologies are created in isolation, rather than as extensions to widely used comprehensive ontologies they miss an opportunity for sharing common meaning. Modern software development, for example in Java or Python, means reusing vast amounts of existing code, such as extensive language libraries and other packages like web servers, databases, device drivers etc. Ontology development needs to follow the same practice to achieve the same efficiency of process as procedural software development.

A second element is the expressivity of definitions. If, as with several of the OAEI ontologies, one must guess at the intended meaning of a term only by its name, then there isn't much chance for shared meaning amongst its users. Each user will just be making a guess. If each term has only a set of binary relationships to other terms then it should still be clear that issues like mutual constraints on values and boundary cases are left unformalized and also at risk of being in conflict among its users. Additionally, without a computational formalization of such constraints, the computer will not be able to test or enforce them. Comments in natural language, no matter how extensive or precise, will not overcome the

[5] http://sigma.ontologyportal.org:8080/sigma/Browse.jsp?lang=EnglishLanguage&flang=SUO-KIF&kb=SUMO&term=FormalMeeting

need for computational definitions, and our experience in this matching effort has been that comments are often not even present, and rarely extensive or precise.

In our work, we were not engaged in correcting the ontologies, since they are part of public datasets. However we consider that a discussion about the problems identified is necessary. Perhaps more robust alignment processes would inherently require modifications in target domain ontologies but also certainly a more detailed formalization. Given the paucity of definitions, we are limited primarily to linguistically-based matches and use of WordNet is a suitable choice for assisting with this sort of match.

5.2 Aligning Conference Ontologies with DOLCE

This section describes the overall methodology we have followed to create the consensual alignment between DOLCE and the domain ontologies. Similarly to SUMO Alignment methodology, this process relies on the existing alignment between DUL and WordNet. We consider these previous alignments as correct so the aim here is to identify the right WordNet synset to the domain concepts. Also, we have reduced the problem to the first-level concepts of the hierarchies from the domain ontologies. This has resulted in 70 first-level domain concepts. For each first level concept of the domain ontology, a top-level specific concept was associated.

We adopt the same set of domain concepts considered by the evaluators to generate correspondences with SUMO. In this way, the senses that better express the domain meaning in WordNet were evaluated and discussed previously. Therefore, we consider the same WordNet correspondence to identify DUL correspondences. Table 5.3 shown an excerpt of the conference-DUL correspondences.

5.3 Existing aligned ontologies

Besides the alignments constructed in the context of this thesis, we adopted as reference alignments some domain ontologies that have been modeled from top-level ontologies. These aligned ontologies can be used as a reference for evaluation of automatic alignment tools.

Table 5.3 – Excerpt of DUL correspondences. Boldface concepts represent the chosen correspondence.

Domain concept	WordNet synsets	DUL concept
cmt#Decision	S1: decision determination conclusion S2: decision determination conclusion S3: decisiveness decision S4: decision S5: decision	S1: **non-physical-collection** S2: activity S3: quality S4: course S5: course
cmt#Preference	S1: predilection preference orientation S2: preference S3: preference predilection taste S4: preference druthers	S1: **cognitive-modal-description** S2: quality S3: cognitive-event S4: cognitive-event
edas#PersonalHistory	S1: history S2: history S3: history account chronicle story S4: history S5: history	S1: temporal-region S2: subject S3: **information-object** S4: temporal-region S5: non-physical-collection
sigkdd#Award	S1: award awarding S2: award accolade honor honour laurels S3: prize award	S1: activity S2: information-object S3: **qualitative-role**
conf#ReviewPreference	S1: predilection preference orientation S2: preference S3: preference penchant predilection taste S4: preference druthers	S1: **cognitive-modal-description** S2: quality S3: cognitive-event S4: cognitive-event

5.3.1 CORA ontology

CORA (IEEE Core Ontology for Robotic and Automation) [54] is an effort of the IEEE Ontologies for Robotics and Automation Working group (ORA). It specifies the main concepts, relations, and axioms of robotics and automation domains. The ontology is composed by four modules (CoraX, Cora, RParts, and POS). CORAX represents concepts and relations commonly found in the ontology that are general to be included in CORA. The module RPARTS includes concepts useful to represent robot parts, and the module POS captures general notions about position and orientation. CORA is aligned to the SUMO top-level ontology. CORA, considering all its modules (CoraX, Cora, RParts, and POS), is composed of 34 first level concepts, from which 29 of them are aligned to SUMO. Table 5.4 shows the classes in CORA which are aligned via a subclass relation to SUMO. A more detailed discussion is available in [54].

5.3.2 SSN (W3C Semantic Sensor Network Ontology)

SSN [23, 24] is an effort of the joint W3C (World Wide Web Consortium) and OGC (Open Geospatial Consortium) Spatial Data on the Web (SDW) Working Group. It describes sensors, devices, observations, measurements and other terms, enabling reasoning of individual sensors and the connection of them. SSN follows a horizontal and vertical modularization architecture by including a lightweight but self-contained core ontology called

Table 5.4 – SUMO to CORA alignments.

corax:DesignObject	subclass of	sumo:abstract
corax:ArtifactSystem	subclass of	sumo:artifact
corax:PhysicalEnvironment	subclass of	sumo:object
corax:Design	subclass of	sumo:proposition
cora:Robot	subclass of	sumo:agent/device
cora:RobotInterface	subclass of	sumo:device
cora:RobotGroup	subclass of	sumo:group
cora:RoboticEnvironment	subclass of	sumo:object
cora:RoboticSystem	subclass of	sumo:artifact
rparts:Actuator	subclass of	sumo:device
rparts:ArmDevice	subclass of	sumo:device
rparts:EndEffector	subclass of	sumo:device
rparts:LocomotionDevice	subclass of	sumo:device
rparts:MechanicalLink	subclass of	sumo:device
rparts:RobotChassis	subclass of	sumo:device
rparts:RoboticBase	subclass of	sumo:device
rparts:WristDevice	subclass of	sumo:device
rparts:ElectronicDevice	subclass of	sumo:electricDevice
rparts:Sensor	subclass of	sumo:measuringDevice
rparts:WheelLocomotionDevice	subclass of	sumo:wheel
corapos:OrientationCoordinateSystem	subclass of	sumo:abstract
corapos:PositionCoordinateSystem	subclass of	sumo:abstract
corapos:OrientationMeasure	subclass of	sumo:physicalQuantity
corapos:PoseMeasure	subclass of	sumo:physicalQuantity
corapos:PositionMesure	subclass of	sumo:physicalQuantity
corapos:OrientationSpatialOperator	subclass of	sumo:function
corapos:OrientationTransformationFn	subclass of	sumo:function
corapos:PositionSpatialOperator	subclass of	sumo:function
corapos:PositionTransformationFn	subclass of	sumo:function

SOSA (Sensor, Observation, Sample, and Actuator) for its elementary classes and properties. Vertical modules build upon each other, hence, they directionally *owl:import* lower level modules. Lower level modules are independent of their higher level modules and logically consistent on their own. In this way, the Dolce-UltraLite Alignment Module imports the SSN Ontology which itself imports SOSA Ontology. Table 5.5 shows classes in SOSA and SSN which are aligned via a subclass relation to DUL. The SSN ontology, considering all its modules (SOSA, SSN) is composed of 18 first level concepts, from those, 8 are aligned to DUL. A more detailed discussion is available in [23, 24].

Table 5.5 – DUL to SOSA and SSN alignments.

sosa:Procedure	subclass of	dul:Method
sosa:Sensor	subclass of	dul:Object
sosa:Observation	subclass of	dul:Event
sosa:Platform	subclass of	dul:Object
ssn:Property	subclass of	dul:Quality
ssn:Stimulus	subclass of	dul:Event
ssn:System	subclass of	dul:Object
ssn:Deployment	subclass of	dul:Event

5.4 Conslusions

This chapter presented the reference alignments considered in our study. The alignments between domain ontologies from the Conference domain with SUMO and DOLCE were built in the context of this work for the purpose of evaluating our approach. Also, we presented two domain ontologies previously built by other research groups. These existing ontologies were based on top-level ontologies [54, 24], which were also adopted as a reference for our evaluations.

We are aware that there are ontologies from other domains, for instance the ontologies available in Bioportal [72] which in principle could be adopted as a reference model. However, these were left out of the scope of our work mainly because they are not aligned with the top-level ontologies under consideration, which are those previously aligned to WordNet. As a continuation of this research, it would be reasonable to consider these other domain ontologies.

Regarding the construction of reference alignments, one of the main issues experts have been faced with concerns the lack of formal definitions associated to terminological layers (comments and labels) that could help to understand the precise semantics of each concept. We recall that although the main motivation for building these alignments were the evaluation of our approach, they can also serve for the evaluation of other systems, as for instance in the context of future OAEI campaigns, and also as a resource for the purpose of semantic bridges in domain ontology matching.

6. EXPERIMENTS

This chapter presents the experiments carried out in order to evaluate our approach. First we describe the material and methods, then we detail how the experiments were conducted, and finally, we discuss the obtained results. As a baseline, we compare our approach with available matching systems participating in the 2017 OAEI campaign. In the end of the chapter, we discuss challenges in the alignment of domain and top-level ontologies, strengths and weaknesses of our developed approach.

6.1 Material and methods

This section presents material and methods adopted in our experiments. First, we present the domain ontologies involved in the process. Then, we describe the resources adopted as an intermediary layer in order to identify correspondences of domain and top-level concepts. Finally, we show systems to be considered as the baseline to compare with our approach.

6.1.1 Ontologies

For our evaluation, we have considered ontologies from three different domains: *(i)* OAEI Conference ontologies, *(ii)* CORA (IEEE Core Ontology for Robotic and Automation), and *(iii)* SSN (W3C Semantic Sensor Network Ontology). As top-level ontologies, we have taken DUL and SUMO.

OAEI Conference ontologies

Seven ontologies from the OAEI Conference data set[1] have been used in our experiments (Cmt, ConfTool, Edas, Ekaw, Iasted, Sigkdd, SofSem). They comprise a total of 70 first-level concepts. These 70 domain concepts were aligned to DUL and SUMO. A description of the manual alignment of the domain and top ontologies was given in Chapter 5.

[1] http://oaei.ontologymatching.org/2018/conference/index.html

CORA

Four modules from Robotics and Automation ontology have been used in our evaluation (CORA, CORAX, RPARTS, and POS). They provide 34 first-level concepts, 29 of them are aligned with SUMO top-level ontology. A full description of the ontology development is available on [54].

SSN (W3C Semantic Sensor Network Ontology)

Two modules from Semantic Sensor Network Ontology have been used (SOSA and SSN). These ontologies provide 18 first-level concepts, where 8 are aligned to DUL. A detailed presentation is available on [23, 24].

Top-level ontologies

Our experiments were conducted considering 2 top-level ontologies, *DUL* and *SUMO*. DUL is composed by 76 concepts and SUMO is composed by 4.558 concepts (counted by Protégé metrics). Chapter 2 details the main concepts and characteristics of each ontology. Table 6.1 shows a summary of the pairs of domain and top ontologies considered in this study.

Table 6.1 – Summary of domain ontologies.

Domain ontology	Modules	First-level concepts	Aligned concepts	Top ontologies
OAEI Conference	7	70	70	DUL and SUMO
CORA	4	34	29	SUMO
SSN	2	18	8	DUL

6.1.2 External resources

Our experiments consider with WordNet as a semantic resource. As presented in Chapter 2, the two adopted top-level ontologies are aligned to WordNet. These previous alignments have been developed by specialists and were adopted as an intermediary layer to identify correspondences of domain and top-level concepts in some experiments. Hence, when we adopt the previous WN-top ontology alignments, if the selected synset is correct, the top-level concept (aligned as super-concept of that synset) is assumed to be a super-concept of the domain concept.

6.1.3 Word embedding models

In addition to the previous WN-top ontology alignments, we conducted experiments considering WE pre-trained models, GloVe [53] and GoogleNews[2]. GloVe is an *unsupervised learning algorithm to obtain vector representations for words*[3]. The adopted trained model uses Wikipedia 2014 and Gigaword5 corpora. It has 6 billions tokens, 400 thousand vocabulary size and neural network dimension of 200. The GoogleNews model is trained on part of Google News dataset (about 100 billion words) via Word2Vec algorithm. The model contains 300-dimensional vectors for 3 million words and phrases.

6.1.4 OAEI tools

Our baseline corresponds to the results of a set of matching tools participating in OAEI 2017, with exception only of those specialised in instance matching (Legato, I-match and njuLink) and one specialised in the bio domain (Yam-bio). The matchers that were tested in our experiment are: ALIN, AML, CroLOM, KEPLER, LogMap, LogMap-Lite, ONTMAT, POMap, SANOM, WikiV3, WikiMatch and XMap. The reader can refer to OAEI papers[4] for a detailed description of them. All tools were run with their default configuration settings.

6.2 Experiments design

The experiments were executed for each domain and their corresponding top-level ontologies, considering different techniques. The experiments are of two different types. The first considers the previous alignment of top ontologies and WordNet. The second is executed without considering previous alignments, that is, the matcher tries to identify the related concept directly with the top ontology.

As mentioned earlier, we run each technique with the domain ontologies against DUL and SUMO. Then we evaluate the generated correspondences considering the reference alignments presented in Chapter 5. In order to evaluate the results, we adopted the same metrics considered in the OAEI campaigns, which are Precision, Recall and F-measure. We further present a qualitative evaluation and an error analysis.

[2]https://code.google.com/archive/p/word2vec/
[3]https://nlp.stanford.edu/projects/glove/
[4]http://www.om2017.ontologymatching.org/#ap

6.2.1 Approaches considering previous WordNet alignments

In this first evaluation, we run classical word sense disambiguation techniques regarding all pairs of ontologies (domain and top). After disambiguating each domain concept, we consider the given previous alignments to WordNet as a way to identify top-level concepts. Table 6.2 shows the number found correspondences for each set of tests.They give an idea of the WordNet coverage for the list of concepts in each domain. As mentioned, for domains which are rather general it would be reasonable to deal with on the basis of this approach.

Table 6.2 – Word sense disambiguation techniques, pair of ontologies and found correspondences considering WordNet previous alignments.

Set of tests	Found correspondences
Conf-DUL	69/70
Conf-SUMO	69/70
SSN-DUL	5/8
Cora-SUMO	26/29

6.2.2 Approaches without previous WordNet alignments

As explained in Chapter 4, we also ran experiments without considering previous alignments with WordNet. In this scenario, we consider the weight of similarity between domain and top concepts on the basis of pre-trained word embedding models. The number of found correspondences is shown on Table 6.3.

Table 6.3 – Word sense disambiguation techniques, pair of ontologies and found correspondences without previous WN alignments.

Set of tests	Found correspondences
Conf-DUL	70/70
Conf-SUMO	70/70
SSN-DUL	8/8
Cora-SUMO	29/29

6.3 Results

We ran our system with the Lesk Similarity (*lesk*), Leveinstein Distance (*LD*), and word embedding models (*WE-GloVe and WE-GNews*) for 16 matching tasks (SSN and DUL, CORA and SUMO, and 7 Conference ontologies with DUL and SUMO). Moreover, we ran the same tasks without previous WordNet alignments, only by considering the concepts similarity based on word embedding models from GloVe and GNews. All alignments generated by our approach are available online[5].

6.3.1 Results considering previous WordNet alignments

Table 6.4 presents our obtained results, considering each pair of ontologies. First column contains the ontology pair, second column corresponds to number of concepts to be aligned to the top-level concepts, the third collumn has the total of domain concepts that were aligned. Last four colums show the correct correspondences found by each disambiguation method (Lesk, Leveinstein Distance, WE-GloVe, and WE-GNews, respectively), according the reference alignments presented in Chapter 5.

Note that only with *Sigkdd* the total of correspondences and found correspondences are different (considering the set of conference domain ontologies). Only one concept was not found in WordNet *(sponzor)*, resulting therefore a different value.

We recall that in the disambiguation step, the selected WordNet synset corresponding to the domain concepts are the same both for DUL and SUMO. So the results are given together in the same line for DUL and SUMO. The seven sub-ontologies for Conference are represented in this table (Cmt, Conf, ConOf, Edas, Ekaw, Iasted, Sigkdd).

Table 6.4 – Tasks and found correct correspondences by each disambiguation method

Task	Total Corresp.	Found corresp.	Lesk	LD	WE-GloVe	WE-GNews
Cmt-DUL/SUMO	8	8	5	3	6	5
Conf.-DUL/SUMO	14	14	11	5	8	10
ConfOf-DUL/SUMO	7	7	6	4	7	6
Edas-DUL/SUMO	16	16	11	6	10	9
Ekaw-DUL/SUMO	6	6	4	1	6	6
Iasted-DUL/SUMO	10	10	7	2	3	4
Sigkdd-DUL/SUMO	9	8	6	5	5	5
SSN-DUL	8	5	0	0	0	0
CORA-SUMO	29	23	9	9	9	9

The generated correspondences have been evaluated in terms of Precision, Recall, and F-measure (Table 6.5).

Table 6.6 groups the results by domain. The for the conference domain, with .72 of F-measure with Lesk. We observe that overall Lesk performs better than Leveinstein, WE-GloVe and WE-GNews. However, looking at the SSN and CORA domain ontologies, the obtained results are lower than for Conference. In fact, for the pair SSN-DUL, we are not be able to find a correct correspondence directly, and for the pair CORA-SUMO, the results

[5]https://github.com/danielasch/top-match

Table 6.5 – Precision, Recall, and F-measure for each disambiguation method.

Task	Lesk			Leveinstein			WE-GloVe			WE-GNews		
	P	R	F-M.	P	R	F-M.	P	R	F-M.	P	R	F-M.
Cmt-DUL/SUMO	0.63	0.63	0.63	0.38	0.38	0.38	0.75	0.75	0.75	0.63	0.63	0.63
Conf.-DUL/SUMO	0.79	0.79	0.79	0.36	0.36	0.36	0.57	0.57	0.57	0.71	0.71	0.71
ConfOf-DUL/SUMO	0.86	0.86	0.86	0.57	0.57	0.57	1.00	1.00	1.00	0.86	0.86	0.86
Edas-DUL/SUMO	0.69	0.69	0.69	0.38	0.38	0.38	0.63	0.63	0.63	0.56	0.56	0.56
Ekaw-DUL/SUMO	0.67	0.67	0.67	0.17	0.17	0.17	1.00	1.00	1.00	1.00	1.00	1.00
Iasted-DUL/SUMO	0.70	0.70	0.70	0.20	0.20	0.20	0.30	0.30	0.30	0.40	0.40	0.40
Sigkdd-DUL/SUMO	0.75	0.67	0.71	0.63	0.56	0.59	0.63	0.56	0.59	0.63	0.56	0.59
SSN-DUL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
CORA-SUMO	0.39	0.31	0.35	0.39	0.31	0.35	0.39	0.31	0.35	0.39	0.31	0.35

were similar, independent of the disambiguation method. Our hypothesis is that concepts from the conference ontology are more general (common sense) than these other domains.

Table 6.6 – Precision, Recall, and F-measure for each disambiguation method, compiled by domain.

Domain	Lesk			Leveinstein			WE-GloVe			WE-GNews		
	P	R	F-M.	P	R	F-M.	P	R	F-M.	P	R	F-M.
Conf.-DUL/SUMO	0.72	0.71	0.72	0.38	0.37	0.37	0.65	0.64	0.65	0.65	0.64	0.65
SSN-DUL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
CORA-SUMO	0.39	0.31	0.35	0.39	0.31	0.35	0.39	0.31	0.35	0.39	0.31	0.35

6.3.2 Results without previous WordNet alignments

In this scenario, we ran the set of 16 matching tasks considering the WE-GloVe and WE-GNews similarity. As presented in Table 6.3 we were able to found correspondences for all domain concepts involved in our experiments, however, when compared to the reference, there is no correct correspondences.

An excerpt of found correspondences is presented in Table 6.7[6]. First column corresponds to the domain concept from *cmt ontology* (conference domain), second column shows the related correspondence with SUMO, and the third column present the concept of SUMO aligned to the domain concept run with WE-GloVe model.

Analysing each output individually, it is possible to observe indirect correspondences, as the correspondence involving the domain concept *Document*. In the reference, it was considered as a subclass of *FactualText*, on the other hand, the generated correspondence considers *Document* as a subclass of *TextDocument*.

In SUMO, *FactualText* is defined as *"The class of Texts that purport to reveal facts about the world. Such texts are often known as information or as non-fiction. Note that something can be an instance of FactualText, even if it is wholly inaccurate. Whether some-*

[6]All generated correspondences are available in https://github.com/danielasch/top-match.

thing is a FactualText is determined by the beliefs of the agent creating the text.", while *TextDocument* is defined as *"Instances of TextDocument are Documents that have at least one part that is an instance of Text."*. Both are subclasses of SUMO concept *Document*, that is defined as *"... Formally, a Document constitutes any ContentBearingObject that is an Artifact conventionally typically intended to be transmitted and assimilated as a meaningful whole. An Article or a Book would be a Document, but a Word or Paragraph typically would not."*

Note that SUMO *Document* is more general and comprise the meaning of a conference document, however, in this context *FactualText* sounds more appropriate to be adopted as a superclass of the domain concept, while *TextDocument*, although an specialization of *Document*, remains with a more generic definition.

We also observe that SUMO concept *NamePart*, was aligned with many domain concepts. For instance, *Conference*, *ProgramCommittee*, *SubjectArea*, and others suppressed of the excerpt. SUMO *NamePart* is a subclass of *ContentBearingObject* that represents *"Any SelfConnectedObject that expresses content. This content may be a Proposition, e.g. when the ContentBearingObject is a Sentence or Text, or it may be a representation of an abstract or physical object, as with an Icon, a Word or a Phrase."*, therefore, analysing the meaning of the concept, it is not a direct correspondence with domain concepts, however, it can be seen as a generalization of them.

In this phase, we do not compute precision, recall and F-measure for generated correspondences, because as mentioned, any found correspondence it was correct, according our reference. However, the experiments shows the challenge of understant the meaning of the concepts to generate correspondences for specialized domains and ontologies with different levels of abstraction.

Table 6.7 – Excerpt of correspondences found by our approach considering the similarity of Word Embedding Models.

Domain Concept	**SUMO Reference**	**SUMO found correspondence**
Bid	Requesting+	OfferingForSale
Conference	FormalMeeting=	NamePart
Decision	Learning+	LegalOpinion
Document	FactualText+	TextDocument
Person	Human=	NamePart
Preference	IntentionalRelation+	VoterAgeRequirement
ProgramCommittee	Commission=	NamePart
SubjectArea	FieldOfStudy=	NamePart

6.3.3 OAEI 2017 matching tools

Only 4 from 12 tested tools (AML, LogMap, LogMapLite, and POMap) were able to find correspondences. They were found for 6 (out of 16) pairs of ontologies, 5 ontologies from conference and CORA. Considering the correspondences found by these tools, 13 domain concepts from conference (out of 70) were aligned. Regarding the number of correspondences, *AML* was able to find 12 correspondences, and 7 of them were correct. *POMap* found 7 correspondences, and 6 were correct. *LogMap* and *LogMapLite* found 6 correspondences respectively, and 5 of them were correct. From the 13 correspondences generated by the matchers, 3 of them were also found by our approach (*confOf:Event* with *DLP:Event*, *edas:ConferenceEvent* with *DLP:Event*, and *iasted:activity* with *DLP:Activity*). Regarding CORA, 1 correspondence was correctly found by *POMap*.

Table 6.8 – Summary of correspondences found by OAEI matching tools.

Domain	Total	Found	AML	LogMap	LogMapLite	POMap
Conference-DUL/SUMO	70	13	12	7	6	6
SSN-DUL	8	0	0	0	0	0
CORA-SUMO	29	1	0	0	0	1

Table 6.8 summarizes the total of found correspondences by matching systems. Note that we have 13 different found correspondences from Conference domain, This value corresponds to the total of different correspondences considering all matching tools.

Table 6.9 shows the results of each matching tool in terms of precision, recall and F-measure. As shown, our approach outperforms all system in terms of Recall and F-measure. Moreover, we compare the results obtained by our approach (named by the adopted disambiguation method). As somehow expected, while the tools perform well in terms of precision, they retrieve a limited number of correspondences, thus, we outperformed them mainly in recall.

Table 6.9 – Precision, Recall, and F-measure for each disambiguation method, comparing with OAEI tools.

Domain	Conference-DUL/SUMO			SSN-DUL			CORA-SUMO		
	P	**R**	**F-M.**	**P**	**R**	**F-M.**	**P**	**R**	**F-M.**
Lesk	0.72	0.71	0.72	0.00	0.00	0.00	0.39	0.31	0.35
Leveinstein	0.38	0.37	0.37	0.00	0.00	0.00	0.39	0.31	0.35
WE-GloVe	0.65	0.64	0.65	0.00	0.00	0.00	0.39	0.31	0.35
WE-GNews	0.65	0.64	0.65	0.00	0.00	0.00	0.39	0.31	0.35
AML	0.92	0.17	0.29	0.00	0.00	0.00	0.00	0.00	0.00
LogMap	0.54	0.10	0.17	0.00	0.00	0.00	0.00	0.00	0.00
LogMapLite	0.46	0.09	0.14	0.00	0.00	0.00	0.00	0.00	0.00
PoMap	0.46	0.09	0.14	0.00	0.00	0.00	1.00	0.03	0.07

6.4 Discussion and Error analysis

In the sections above, we presented and discussed the experiments. In addition, we compared our methods with correspondences found by matching tools participating of

OAEI campaigns. Although our approach was able to found a high number of correspondences for the three domains, in some cases, the generated correspondences were wrong. First, the context of concepts that we adopted seems not to be enough to disambiguate the sense of the domain concept (Conference domain ontologies are not equipped of comments and labels). This could be improved by enriching the terminological layer. Second, we observed that the adopted word embedding models can contribute to the disambiguation step, however, they were not better than the classical approaches. An alternative test would be to use domain-specific embedding models since we have used models based on general corpora. Third, the word sense disambiguation here is still based on the overlapping of words, and word sense disambiguation techniques could be adopted as an alternative way, by retrieving the similarity os terms.

Observing all presented results, it is interesting to analyse how methods can be adopted in a complementary way, for instance, Table 6.6 shows that *Lesk* obtained better results for the conference domain when compared to the other adopted methods. On the other hand, *Leveinstein* presents the worst ratio of precision, recall, and f-measure. But, if we consider a qualitative analysis, the adopted methods are able to find different correspondences, i.e. the domain concept *Bid*, that corresponds to four senses for in WordNet:

1. bid tender | a formal proposal to buy at a specified price.

2. bid play | an attempt to get something; "they made a futile play for power"; "he made a bid to gain attention".

3. bid bidding | (bridge) the number of tricks a bridge player is willing to contract to make.

4. command bid bidding dictation | an authoritative direction or instruction to do something.

Each sense is aligned with a different top-level concept. In the domain, the sense that better express the meaning of the concept is the first one, which is aligned with *Requesting* in SUMO. Looking for the output of the methods, (i) *Lesk* align with *Ordering*, (ii) *Leveinstein* align with *Committing*, and both (iii) word embedding models align with *Requesting*. Table 6.10 presents examples of correspondences that is interesting to observe in a qualitative way.

Similarly, the concept (i) *Conference* was aligned correctly with *Lesk* and *Leveinstein*, but not with Word Embedding similarity. (ii) *Document* was aligned correctly with three methods, (iii) regarding the concept *Person*, it is interesting to observe that it was aligned with two top concepts (*Human* and *Word*), by the same method, in different domain ontologies. If we go back to WordNet, *Person* has three senses, first and second related to the top concept *Human*, and third corresponds to a grammatical category that is aligned with the top concept *Word*:

Table 6.10 – Examples of Found Correspondences.

Domain Concept	**SUMO Reference**	**Lesk**	**Leveinstein**	**WE-GloVe**	**WE-GNews**
cmt#Conference	Formal Meeting	Formal Meeting	Formal Meeting	Communication	Communication
cmt#Document	FactualText	FactualText	ComputerFile	FactualText	FactualText
cmt#Person	Human	Human	Word	Human	Human
sigkdd#Person	Human	Word	Word	Human	Human
ConfOf#Contribution	Text	Entity	Entity	Text	Text
edas#Call	Requesting	Subjective Assessment Attribute	Meeting	Subjective Assessment Attribute	Requesting
edas#ReviewForm	Text	Word	Human	Text	Proposition
iasted#Activity	Intentional Process	Intentional Process	Process	Intentional Process	Intentional Process
sigkdd#Place	Geographic Area	Position	Region	Region	Position
edas#Place	Geographic Area	LandArea	Region	Subjective Assessment Attribute	LandArea
rparts:RoboticBase	device	Military Installation	Military Installation	Military Installation	Military Installation

1. person individual someone somebody mortal soul | a human being; "there was too much for one person to do"

2. person | a human body (usually including the clothing); "a weapon was hidden on his person"

3. person | a grammatical category used in the classification of pronouns, possessive determiners, and verb forms according to whether they indicate the speaker, the addressee, or a third party; "stop talking about yourself in the third person"

This example also highlights the importance of the context in the disambiguation process, because we can observe that the same method (*Lesk*) found different correspondence in a same domain, for different ontologies.

In addition, for the concept Contribution, Lesk and Leveinstein found the same correspondence, while WE models also find a same correspondence. In this case, the correct correspondence was found by WE models. On the other hand, there are correspondences found correctly only by one method, as the case of *Call* and *ReviewForm*. For the case of *ReviewForm* each method has found a different correspondence.

Place is a concept presents in more domain ontologies, it should be aligned with SUMO *GeographicArea*, however, the found correspondence was wrong to all methods and in both cases. With *Lesk*, and *Word Embedding* models, the correspondence was different in each task, but neither was correct, according to the reference. On the other hand, Looking the meaning of the concepts, *GeograficArea* is defined in SUMO as a subclass of *Region* and *LandArea* is a subclass of *GeograficArea*. Hence, some correspondences could be inferred or validated via transitivity relation.

Regarding the concept *RoboticBase*, all methods found as correspondence the concept *MilitaryInstalation*. Considering the meaning of the domain concept, *Device* corresponds to the best match, once is defined in SUMO as *"A Device is an Artifact whose purpose is to serve as an instrument in a specific subclass of Process."*, while *MilitaryInstalation* represents *" A StationaryArtifact consisting of grounds and Buildings that is intended to be used by a MilitaryOrganization."*.

In all analysed incorrect correspondences, we observe the lack of context to help in the disambiguation process. For instance, if we analyze the concept *Conference* in the domain ontology *cmt*, the context of the concept is composed only by the name of the concept, because there is no sub, sup, or annotation related to them. Hence, although there are three senses available in WordNet, there is only one term to overlap to retrieve the ratio of similarity.

6.5 Conclusions

In this chapter we detailed the experiments executed in order to evaluate our approach. We discussed correspondences generated by our different methods and compare them with matching systems participating of OAEI campaigns. Our results highlights the importance of the context of the concept to disambiguate the meaning of the concept. Moreover, we observed that a combination of methods can be a way to improve the performance of our approach. On the other hand, if we compare our results regarding available matching tools, it is evident the improvement of recall and consequently, F-measure. Regarding the methods which do not adopt WordNet as an intermediary layer, we observed that the generated correspondences are mainly wrong. In this way, it is important to extend the experiments to evaluate other mechanisms in order to generate correspondences with a higher precision and recall.

7. CONCLUSION

In this work we studied the important but so far quite unattended issue of automatic aligning domain and top ontologies. Since this is a very difficult task with no previous available resources, as a first approach, and one that was not taken into account before, we have investigated the use of WordNet as an auxiliary resource to achieve the alignments. Considering the investigated domains, which are rather general than specialized, WordNet has been proven a useful external resource in the task of aligning top-level and domain ontologies.

To match the domain concepts with a WordNet synset, we combined classical WordNet measures with other distributional semantics approaches. We adopted the classical word sense disambiguation methods such as Lesk and Leveinstein distance to identify the WordNet sense that better express the meaning of the domain concept. In addition, we combined the similarity of word embedding models to retrieve the ratio of similarity of the context of the concept and the context of each WordNet sense as an alternative to the classical methods.

Regarding the obtained results, our approach outperformed OAEI matching systems in terms of recall. The best result was achieved with *Lesk*, for conference domain adopting previous WordNet alignments as an intermediary layer to identify correspondences with 0.72 of F-measure. Word embedding models obtained similar results, 0.65 of F-measure. Although Leveinstein did not performed so well, it presented better F-measure then OAEI systems. Moreover, it is possible to observe that all methods and systems were better for the Conference domain.

But there is indeed an important usual limitation regarding the use of the lexical database WordNet, since it presents mostly general terms and also it does not contain compound terms. For ontologies of very specialized domains, their use could be inefficient. Also, regarding the previous alignment involving WordNet and top-level ontologies we are limited to the version of WordNet and top ontologies used in the alignments. Besides these limitations, considering the lack of previous works dealing with this problem we consider that this first original proposal and its analysis have shown a practical useful way to improve the process of aligning domain and top ontologies, in domains that are not very specialized and within certain limits.

We analysed available domain ontologies aligning tools, which were not developed for the purpose of considering top ontologies, but which were the only practical automatic aligning tools available. As some intend to cover also subsumption relations, it could be the case that they could perform in this special case of aligning domain and top ontologies.

The results have shown that align domain and top-level ontologies remains a hard work task and requires deep knowledge of the semantics of the information encoded in the ontologies, however, our approach starts a discussion of how we can work with ontologies from different levels of abstraction in computational application and real world scenarios.

Hence, matching systems need to be improved to include more abstract and philosophical semantic relations and semiotic matching, to take advantage of structural features of the ontologies and axioms in order to better compare their formal definitions, and to take advantage of background knowledge from external resources targeting subsumption and other semantic relations. These have to be combined with logical reasoning techniques for guarantee the consistency of the generated alignments. The current approaches have to revised to better deal with the specificities of matching with top ontologies.

Last but not least, another contribution of this work is the reference alignment involving the OAEI Conference dataset and DOLCE and SUMO ontologies, that was developed in the context of this thesis. On the basis of this new reference alignment it would be possible to propose a first OAEI track for the task of domain and top ontology alignment.

7.1 Future Work

As future work we plan to *(i)* consider ways of adding context to the terms; *(ii)* run experiments in order to combine methods to improve precision and recall of the generated alignment. Besides, *(iii)* it is also important to focus on multi word expressions, mainly to work with specialized domains. We also intend *(iv)* to extend our experiments to cover other domains; *(v)* develop reference alignments to be adopted as a gold standard for future evaluations; and *(vi)* explore alternative ways to align domain and top-level ontologies without the dependency of WordNet.

7.2 Publications

During the progress on the thesis, the research and some results were presented and discussed in the following papers:

- SCHMIDT, Daniela; TROJAHN, Cassia; VIEIRA, Renata. Analysing Top-level and Domain Ontology Alignments from Matching Systems. In: Proceedings of the Eleventh International Workshop on Ontology Matching - OM, 2016, pp. 13-24.

- SCHMIDT, Daniela; TROJAHN, Cassia; VIEIRA, Renata; Kamel, Mouna. Validating Top-level and Domain Ontology Alignments from Matching Systems. In: Seminar on Ontology Research in Brazil - ONTOBRAS, 2016, pp. 119-130.

- SCHMIDT, Daniela; BASSO, Rafael; TROJAHN, Cassia; VIEIRA, Renata. Matching Domain and Top-level Ontologies via OntoWordNET. In: Proceedings of the Twelfth International Workshop on Ontology Matching - OM, 2017, pp. 225-226.

- BASSO, Rafael; SCHMIDT, Daniela; TROJAHN, Cassia; VIEIRA, Renata. Top-level and Domain Ontologies Alignment via WordNet. In: Seminar on Ontology Research in Brazil - ONTOBRAS, 2017, pp.9-20.

- SCHMIDT, Daniela; BASSO, Rafael; TROJAHN, Cassia; VIEIRA, Renata. Matching Domain and Top-level ontologies exploring Word Sense Disambiguation and Word Embedding. In: Proceedings of the Thirteenth International Workshop on Ontology Matching - OM, 2018, pp. 1-12 (Best Workshop Paper).

- SCHMIDT, Daniela; BASSO, Rafael; TROJAHN, Cassia; VIEIRA, Renata. Matching Domain and Top-level ontologies exploring Word Sense Disambiguation and Word Embedding. In: Emerging Topics in Semantic Technologies. E. Demidova, A.J. Zaveri, E. Simperl (Eds.). ISBN: 978-3-89838-736-1. (c) 2018 AKA Verlag Berlin.

- KAMEL, Mouna; SCHMIDT, Daniela; TROJAHN, Cassia; VIEIRA, Renata. Exploring BabelNet for generating subsumption. In: Proceedings of the Thirteenth International Workshop on Ontology Matching - OM, 2018, pp. 1-2.

- SCHMIDT, Daniela; PEASE, Adam; TROJAHN, Cassia; VIEIRA, Renata. Aligning Conference Ontologies with SUMO: A Report on Manual Alignment via WordNet. In: Proceedings of the Joint Ontology Workshops (Workshop on Foundational Ontology FOUST III). 2019, pp. 1-10.

- KAMEL, Mouna; SCHMIDT, Daniela; TROJAHN, Cassia; VIEIRA, Renata. Hypernym relation extraction for establishing subsumptions: preliminary results on matching foundational ontologies . In: Proceedings of the Fourteenth International Workshop on Ontology Matching - OM, 2019, pp. 1-5.

- SCHMIDT, Daniela; TROJAHN, Cassia; VIEIRA, Renata. Matching BFO, DOLCE, GFO and SUMO: an Evaluation of OAEI 2018 Matching Systems. In: Seminar on Ontology Research in Brazil - ONTOBRAS, 2019, pp.1-12.

- SCHMIDT, Daniela; DAL BOSCO, Avner; QUARESMA, Paulo; TROJAHN, Cassia; VIEIRA, Renata. Aligning IATE criminal terminology to SUMO. In: Proceedings of the International Conference on Computational Processing of Portuguese - PROPOR, 2020, pp. 1-11.

- SCHMIDT, Daniela; DAL BOSCO, Avner; QUARESMA, Paulo; TROJAHN, Cassia; VIEIRA, Renata. Evaluating the alignment of multi domain terminologies to SUMO.

In: 12th Edition of its Language Resources and Evaluation Conference - LREC, 2020. [Submitted].

REFERENCES

[1] Brodaric, B.; Probst, F. “DOLCE ROCKS: Integrating Geoscience Ontologies with DOLCE”. In: Proceedings of the Semantic Scientific Knowledge Integration Conference, 2008, pp. 3–8.

[2] Cheatham, M.; Hitzler, P. “Conference v2.0: An uncertain version of the OAEI conference benchmark”. In: Proceedings of the International Semantic Web Conference, 2014, pp. 33–48.

[3] Damova, M.; Kiryakov, A.; Simov, K. I.; Petrov, S. “Mapping the central LOD ontologies to PROTON Upper-Level Ontology”. In: Proceedings of the Workshop on Ontology Matching, 2010, pp. 61–72.

[4] David, J.; Euzenat, J.; Scharffe, F.; dos Santos, C. T. “The Alignment API 4.0”, *Semantic Web Journal*, vol. 2–1, January 2011, pp. 3–10.

[5] Euzenat, J.; Shvaiko, P. “Ontology Matching”. Springer, 2013, 511p.

[6] Fellbaum, C. “WordNet: An Electronic Lexical Database”. MIT Press, 1998, 449p.

[7] Gangemi, A. “DOLCE+DnS Ultralite (DUL)”. Source: http://www.ontologydesignpatterns.org/ont/dul/DUL.owl, Dec 2019.

[8] Gangemi, A.; Guarino, N.; Masolo, C.; Oltramari, A.; Schneider, L. “Sweetening Ontologies with DOLCE”. In: Proceedings of the Knowledge Engineering and Knowledge Management Conference, 2002, pp. 166–181.

[9] Gangemi, A.; Navigli, R.; Velardi, P. “The OntoWordNet Project: Extension and Axiomatization of Conceptual Relations in WordNet”. In: Proceedings of the Move to Meaningful Internet Systems Conference, 2003, pp. 820–838.

[10] Gangemi, A.; Nuzzolese, A. G.; Presutti, V.; Draicchio, F.; Musetti, A.; Ciancarini, P. “Automatic Typing of DBpedia Entities”. In: Proceedings of the Semantic Web Conference, 2012, pp. 65–81.

[11] Gómez-Pérez, A.; Fernández-López, M.; Corcho, O. “Ontological Engineering: With Examples from the Areas of Knowledge Management, e-Commerce and the Semantic Web”. Springer, 2007, 415p.

[12] Grenon, P. “BFO in a Nutshell: A Bi-categorial Axiomatization of BFO and Comparison with DOLCE”, Technical report, Institute for Formal Ontology and Medical Information Science, IFOMIS, 2003, 66p.

[13] Grenon, P.; Smith, B.; Goldberg, L. “Biodynamic Ontology: Applying BFO in the Biomedical Domain”. In: Proceedings of the Studies in Health Technology and Informatics Conference, 2004, pp. 20–38.

[14] Gruber, T. “A Translation Approach to Portable Ontology Specifications”, *Knowledge Acquisition Journal*, vol. 5–2, June 1993, pp. 199–220.

[15] Guarino, N. “Some organizing principles for a Unified Top-Level ontology”. In: Proceedings of the AAAI Spring Symposium on Ontological Engineering, 1997, pp. 1–17.

[16] Guarino, N. “Formal Ontology in Information Systems”. IOS Press, 1998, 347p.

[17] Guarino, N.; Welty, C. A. “An Overview of OntoClean”. In: *Handbook on Ontologies*, Staab, S.; Studer, R. (Editors), Springer Berlin Heidelberg, 2004, chap. 8, pp. 151–171.

[18] Guha, V.; Lenat, D. “Cyc: A Midterm Report”, *AI Magazine*, vol. 11–3, September 1990, pp. 32.

[19] Guizzardi, G. “Ontological Foundations for Structural Conceptual Models”, Ph.D. Thesis, University of Twente, Enschede, The Netherlands, Enschede, 2005, 416p.

[20] Guizzardi, G. “The Role of Foundational Ontologies for Conceptual Modeling and Domain Ontology Representation”. In: Proceedings of the International Baltic Conference on Databases and Information Systems, 2006, pp. 17–25.

[21] Guizzardi, G.; Wagner, G. “Using the Unified Foundational Ontology (UFO) as a Foundation for General Conceptual Modeling Languages”. In: *Theory and Applications of Ontology: Computer Applications*, Poli, R.; Healy, M.; Kameas, A. (Editors), Springer, 2010, chap. 8, pp. 175–196.

[22] Guizzardi, G.; Wagner, G.; Almeida, J. P. A.; Guizzardi, R. S. “Towards Ontological Foundations for Conceptual Modeling: The Unified Foundational Ontology (UFO) Story”, *Applied Ontology Journal*, vol. 10–3-4, June 2015, pp. 259–271.

[23] Haller, A.; Janowicz, K.; Cox, S.; Phuoc, D. L.; Taylor, K.; Lefrançois, M. “Semantic Sensor Network Ontology”. Source: https://www.w3.org/TR/vocab-ssn/, Oct 2019.

[24] Haller, A.; Janowicz, K.; Cox, S. J.; Lefrançois, M.; Taylor, K.; Le Phuoc, D.; Lieberman, J.; García-Castro, R.; Atkinson, R.; Stadler, C. “The Modular SSN Ontology: A Joint W3C and OGC Standard Specifying the Semantics of Sensors, Observations, Sampling, and Actuation”, *Semantic Web Journal*, vol. 10–1, December 2018, pp. 9–32.

[25] Herre, H.; Heller, B.; Burek, P.; Hoehndorf, R.; Loebe, F.; Michalek, H. “General Formal Ontology (GFO): A Foundational Ontology Integrating Objects and Processes”, Technical report, Research Group Ontologies in Medicine, University of Leipzig, 2006, 90p.

[26] Horridge, M.; Bechhofer, S. “The OWL API: A Java API for OWL Ontologies”, *Semantic Web Journal*, vol. 2–1, January 2011, pp. 11–21.

[27] Jain, P.; Yeh, P.; Verma, K.; Vasquez, R.; Damova, M.; Hitzler, P.; Sheth, A. “Contextual Ontology Alignment of LOD with an Upper Ontology: A Case Study with PROTON”. In: Proceedings of the Extended Semantic Web Conference, 2011, pp. 80–92.

[28] Jurafsky, D.; Martin, J. H. “Speech and Language Processing”. Prentice Hall, 2009, 988p.

[29] Kalfoglou, Y.; Schorlemmer, M. “Ontology Mapping: the state of the art”, *Knowledge Engineering Review*, vol. 18–1, January 2003, pp. 1–31.

[30] Keet, C. “The Use of Foundational Ontologies in Ontology Development: An Empirical Assessment”. In: Proceedings of the Extended Semantic Web Conference, 2011, pp. 321–335.

[31] Khan, Z.; Keet, C. “Addressing Issues in Foundational Ontology Mediation”. In: Proceedings of the Knowledge Engineering and Ontology Development Conference, 2013, pp. 5–16.

[32] Khan, Z.; Keet, C. “The Foundational Ontology Library ROMULUS”. In: Proceedings of the Model and Data Engineering Conference, 2013, pp. 200–211.

[33] Khan, Z.; Keet, C. “Feasibility of Automated Foundational Ontology Interchangeability”. In: Proceedings of the Knowledge Engineering and Knowledge Management Conference, 2014, pp. 225–237.

[34] Leao, F.; Revoredo, K.; Baiao, F. “Extending WordNet with UFO Foundational Ontology”, *Journal of Web Semantics*, vol. 57–1, August 2019, pp. 1–17.

[35] Lesk, M. “Automatic Sense Disambiguation Using Machine Readable Dictionaries: How to Tell a Pine Cone from an Ice Cream Cone”. In: Proceedings of the Systems Documentation Conference, 1986, pp. 24–26.

[36] Levenshtein, V. I. “Binary Codes Capable of Correcting Deletions, Insertions, and Reversals”, *Cybernetics and Control Theory Journal*, vol. 10–8, February 1966, pp. 707–710.

[37] Li, J. “LOM: A Lexicon-based Ontology Mapping Tool”. In: Proceedings of the Performance Metrics for Intelligent Systems Workshop, 2004, pp. 1–5.

[38] Mascardi, V.; Cordì, V.; Rosso, P. “A Comparison of Upper Ontologies”. In: Proceedings of the Agents and Industry Conference, 2007, pp. 55–64.

[39] Mascardi, V.; Locoro, A.; Rosso, P. “Automatic Ontology Matching via Upper Ontologies: A Systematic Evaluation”, *Transactions on Knowledge and Data Engineering Journal*, vol. 22–5, July 2009, pp. 609–623.

[40] Meilicke, C.; Garcia-Castro, R.; Freitas, F.; van Hage, W. R.; Montiel-Ponsoda, E.; de Azevedo, R. R.; Stuckenschmidt, H.; Sváb-Zamazal, O.; Svátek, V.; Tamilin, A.; Trojahn, C.; Wang, S. “MultiFarm: A Benchmark for Multilingual Matching”, *Journal of Web Semantics*, vol. 15–1, September 2012, pp. 62–68.

[41] Mika, P.; Oberle, D.; Gangemi, A.; Sabou, M. “Foundations for Service Ontologies: Aligning OWL-S to DOLCE”. In: Proceedings of the Conference on World Wide Web, 2004, pp. 563–572.

[42] Mikolov, T.; Chen, K.; Corrado, G.; Dean, J. “Efficient Estimation of Word Representations in Vector Space”. In: Proceedings of the Learning Representations Conference, 2013, pp. 1–12.

[43] Miller, G. A. “WordNet: A Lexical Database for English”, *Communications of the ACM*, vol. 38–11, November 1995, pp. 39–41.

[44] Muñoz, L.; Grüninger, M. “Verifying and Mapping the Mereotopology of Upper-Level Ontologies”. In: Proceedings of the Knowledge Discovery, Knowledge Engineering and Knowledge Management Conference, 2016, pp. 31–42.

[45] Niles, I.; Pease, A. “Towards a Standard Upper Ontology”. In: Proceedings of the Formal Ontology in Information Systems Conference, 2001, pp. 2–9.

[46] Niles, I.; Pease, A. “Linking Lexicons and Ontologies: Mapping WordNet to the Suggested Upper Merged Ontology”. In: Proceedings of the Information and Knowledge Engineering Conference, 2003, pp. 412–416.

[47] OAEI. “Ontology Alignment Evaluation Initiative”. Source: http://oaei.ontologymatching.org/, Nov 2019.

[48] Oberle, D.; Ankolekar, A.; Hitzler, P.; Cimiano, P.; Sintek, M.; Kiesel, M.; Mougouie, B.; Baumann, S.; Vembu, S.; Romanelli, M.; Buitelaar, P.; Engel, R.; Sonntag, D.; Reithinger, N.; Loos, B.; Zorn, H.; Micelli, V.; Porzel, R.; Schmidt, C.; Weiten, M.; Burkhardt, F.; Zhou, J. “DOLCE Ergo SUMO: On Foundational and Domain Models in the SmartWeb Integrated Ontology (SWIntO)”, *Journal of Web Semantics*, vol. 5–3, September 2007, pp. 156–174.

[49] Padilha, N.; Baião, F.; Revoredo, K. “Alignment Patterns based on Unified Foundational Ontology”. In: Proceedings of the Brazilian Ontology Research Seminar, 2012, pp. 48–59.

[50] Pease, A. “Formal Representation of Concepts: The Suggested Upper Merged Ontology and its Use in Linguistics”. In: *Ontolinguistics: How Ontological Status Shapes the Linguistic Coding of Concepts*, Schalley, A.; Zaefferer, D. (Editors), Vorbereitung Berlin, 2006, chap. 2, pp. 103–114.

[51] Pease, A. “Ontology: A Practical Guide”. Articulate Software Press, 2011, 303p.

[52] Pease, A.; Benzmüller, C. “Sigma: An Integrated Development Environment for Logical Theories”, *AI Communications*, vol. 26–1, January 2013, pp. 79–97.

[53] Pennington, J.; Socher, R.; Manning, C. D. “GloVe: Global Vectors for Word Representation”. In: Proceedings of the Empirical Methods in Natural Language Processing Conference, 2014, pp. 1532–1543.

[54] Prestes, E.; Carbonera, J. L.; Fiorini, S. R.; Jorge, V.; Abel, M.; Madhavan, R.; Locoro, A.; Goncalves, P.; Barreto, M.; Habib, M.; Chibani, A.; Gérard, S.; Amirat, Y.; Schlenoff, C. “Towards a Core Ontology for Robotics and Automation”, *Robotics and Autonomous Systems Journal*, vol. 61–11, November 2013, pp. 1193 – 1204.

[55] Reed, S.; Lenat, D. “Mapping Ontologies into Cyc”. In: Proceedings of the Ontologies For The Semantic Web Conference, 2002, pp. 1–6.

[56] Schadd, F. C.; Roos, N. “Coupling of Wordnet Entries for Ontology Mapping Using Virtual Documents”. In: Proceedings of the Workshop on Ontology Matching, 2012, pp. 25–36.

[57] Schmidt, D.; Basso, R.; Trojahn, C.; Vieira, R. “Matching Domain and Top-level ontologies exploring Word Sense Disambiguation and Word Embedding”, *Emerging Topics in Semantic Technologies Journal*, vol. 36–1, January 2018, pp. 27–38.

[58] Schmidt, D.; Trojahn, C.; Vieira, R. “Analysing Top-level and Domain Ontology Alignments from Matching Systems”. In: Proceedings of the Workshop on Ontology Matching, 2016, pp. 1–12.

[59] Semy, S.; Pulvermacher, M.; Obrst, L. “Toward the Use of an Upper Ontology for U.S. Government and U.S. Military Domains: An evaluation”, Technical report, The MITRE Corporation, 2004, 42p.

[60] Seppälä, S. “Mapping WordNet to Basic Formal Ontology using the KYOTO ontology”. In: Proceedings of the Biomedical Ontology Conference, 2015, pp. 1–2.

[61] Seppälä, S.; Smith, B.; Ceusters, W. “Applying the Realism-Based Ontology-Versioning Method for Tracking Changes in the Basic Formal Ontology”. In: Proceedings of the Formal Ontology in Information Systems Conference, 2014, pp. 227–240.

[62] Seyed, A. “BFO/DOLCE Primitive Relation Comparison”. In: Proceedings of the International Conference on Biomedical Ontology, 2009, pp. 1–2.

[63] Silva, V.; Campos, M.; Silva, J.; Cavalcanti, M. “An Approach for the Alignment of Biomedical Ontologies based on Foundational Ontologies”, *Information and Data Management Journal*, vol. 2–3, January 2011, pp. 557–572.

[64] Silva, V. S.; Freitas, A.; Handschuh, S. “Word Tagging with Foundational Ontology Classes: Extending the WordNet-DOLCE Mapping to Verbs”. In: *Knowledge Engineering and Knowledge Management*, Blomqvist, E.; Ciancarini, P.; Poggi, F.; Vitali, F. (Editors), Springer, 2016, chap. 38, pp. 593–605.

[65] Sowa, J. F. “Knowledge Representation: Logical, Philosophical and Computational Foundations”. Cengage Learning, 1999, 512p.

[66] Stevens, R.; Lord, P.; Malone, J.; Matentzoglu, N. “Measuring Expert Performance at Manually Classifying Domain Entities under Upper Ontology Classes”, *Journal of Web Semantics*, vol. 57–1, August 2019, pp. 1–13.

[67] Swartout, B.; Ramesh, P.; Knight, K.; Russ, T. “Toward Distributed use of Large-Scale Ontologies”. In: Proceedings of the Ontological Engineering Workshop, 1997, pp. 138–148.

[68] Temal, L.; Rosier, A.; Dameron, O.; Burgun, A. “Mapping BFO and DOLCE”. In: Proceedings of the Medical Informatics Conference, 2010, pp. 1065–1069.

[69] Terziev, I.; Kiryakov, A.; Manov, D. “Base Upper-level Ontology (BULO) Guidance”, Technical report, Semantically-Enabled Knowledge Technologies, SEKT Consortium, 2005, 70p.

[70] Thiéblin, É.; Haemmerlé, O.; Hernandez, N.; Trojahn, C. “Task-Oriented Complex Ontology Alignment: Two Alignment Evaluation Sets”. In: Proceedings of the Extended Semantic Web Conference, 2018, pp. 655–670.

[71] Wang, P. “Lily Results on SEALS Platform for OAEI 2011”. In: Proceedings of the Workshop on Ontology Matching, 2011, pp. 156–162.

[72] Whetzel, P. L.; Noy, N. F.; Shah, N. H.; Alexander, P. R.; Nyulas, C.; Tudorache, T.; Musen, M. A. “BioPortal: Enhanced functionality via new Web Services from the National Center for Biomedical Ontology to access and use Ontologies in Software Applications”, *Nucleic Acids Research Journal*, vol. 39–2, July 2011, pp. 541–545.

[73] Yatskevich, M.; Giunchiglia, F. "Element Level Semantic Matching using WordNet". In: Proceedings of the Meaning Coordination and Negotiation Workshop, 2004, pp. 37–48.

[74] Zamazal, O.; Svátek, V. "The Ten-Year OntoFarm and its Fertilization within the Onto-Sphere", *Journal of Web Semantics*, vol. 43–1, March 2017, pp. 46–53.

Pontifícia Universidade Católica do Rio Grande do Sul
Pró-Reitoria de Graduação
Av. Ipiranga, 6681 - Prédio 1 - 3º. andar
Porto Alegre - RS - Brasil
Fone: (51) 3320-3500 - Fax: (51) 3339-1564
E-mail: prograd@pucrs.br
Site: www.pucrs.br