

Fast developing of a Natural Language Interface for a Portuguese Wordnet: Leveraging on Sentence Embeddings

Hugo Gonalo Oliveira

CISUC, Dept. Informatics Engineering
Universidade de Coimbra, Portugal
hroliv@dei.uc.pt

Alexandre Rademaker

IBM Research and EMAP/FGV
Rio de Janeiro, Brazil
arademaker@gmail.com

Abstract

We describe how a natural language interface can be developed for a wordnet with a small set of handcrafted templates, leveraging on sentence embeddings. The proposed approach does not use rules for parsing natural language queries but experiments showed that the embeddings model is tolerant enough for correctly predicting relation types that do not match known patterns exactly. It was tested with OpenWordNet-PT, for which this method may provide an alternative interface, with benefits also on the curation process.

1 Introduction

A natural way of interacting with computational systems or knowledge bases is to use the same language we use for interacting with other humans. However, due to all the complex phenomena of natural language, most systems rely on browsing, keyword-based search interfaces or their combination. This is simpler at the technical level and avoids having to deal with Natural Language Understanding issues. The previous phenomena include ambiguity and language variability and are the reason why matching natural language with formal queries is not a trivial task. To overcome this challenge, we investigate how a model of sentence similarity can be exploited by a natural language interface (NLI) for a wordnet. Our approach is tested in OpenWordNet-PT (de Paiva et al., 2012) (OWN-PT), probably the most active Portuguese wordnet (de Paiva et al., 2016b).

The development of this system, dubbed NELIO, requires only a small set of handcrafted templates for each query to be covered. Instantiating those templates with arguments from OWN-PT results in a large set of sentences, used for training a doc2vec (Le and Mikolov, 2014) model. The latter is a variation of the popular word2vec (Mikolov et al., 2013) but, besides learning dense vector representations of words, it

learns a representation for documents (sentences, in our case), based on the words used and on a document label. Such a model can be used e.g., for predicting the most suitable label for an unseen document. In this work, we rely on the trained doc2vec model for predicting the relation type that a natural language query is asking for. We then use this information for querying OWN-PT and retrieving suitable answers. This process is fast enough and avoids writing a set of rules for parsing natural language queries. Besides providing a more natural way of interacting with OWN-PT, NELIO turns out to be an alternative way of exploring OWN-PT and reveal flaws that, otherwise, would not be easy to spot.

The remainder of this paper briefly overviews OWN-PT, describes the development of NELIO, reports on performed experiments, including a systematic evaluation of the model in this context, and, before concluding, overviews related work.

2 OpenWordNet-PT

OpenWordNet-PT (OWN-PT) is an ongoing project to build a wordnet for Portuguese. It is aligned with Princeton WordNet (Fellbaum, 1998) (PWN), but still has about half of its size. So far, only partial evaluations of its coverage were performed, namely of verbs (de Paiva et al., 2016a, 2014), nouns (Rademaker et al., 2014), and (gentilic) adjectives (Real et al., 2016).

OWN-PT is freely available in RDF/OWL. Its data can be retrieved via a SPARQL endpoint, but it can also be explored through its own web interface¹ or through the interface of the Open Multilingual WordNet (Bond and Foster, 2013). As previously suggested (Real et al., 2015), a visual interface helps to discover interesting issues to work on. The research presented here is related to lessons previously learned.

¹<http://openwordnet-pt.org>

3 System Development

NELIO interprets questions, in Portuguese, that ask for concepts, lexicalised as y , which are related in some way to another concept lexicalised as x , mentioned in the question. This section describes the steps for developing its current version.

3.1 Question Templates

To enable the generation of prototypical questions, a small set of templates for each covered relation was handcrafted by the first author of this paper. Such templates generalise possible ways of asking the desired questions in Portuguese. All templates currently used (between 3 and 10 per relation) are revealed in table 1, grouped according to the target relation. Most semantic relations in OWN-PT are covered. Yet, due to their different scope, lexical relations were left out of this set.

3.2 Model Training

The generation of prototypical questions results from filling the templates, automatically, with real examples from OWN-PT. Those questions were used to train a doc2vec (Le and Mikolov, 2014) model, with the name of the target semantic relation set as their label. Examples of generated questions include:

<i>(hyponymOf)</i>	<i>que formas há de correr?</i>
<i>(hypernymOf)</i>	<i>qual é o hiperónimo de maçã?</i>
<i>(memberHol..Of)</i>	<i>quais os membros de Liga Árabe?</i>
<i>(substanceHol..Of)</i>	<i>de que é feito molho de soja?</i>
<i>(partHolonymOf)</i>	<i>que partes tem Portugal?</i>
<i>(partMeronym)</i>	<i>de que faz parte Breslávia?</i>
<i>(antonymOf)</i>	<i>qual é o contrário de líquido?</i>
<i>(x_causes)</i>	<i>qual é o efeito de ferir?</i>
<i>(entails)</i>	<i>o que implica migrar?</i>

The learned model can be exploited in a classification task. More precisely, given a fragment of text, it can be used for predicting the appropriate label. Once predicted, the label is used together with the relation argument that appears on the question (x) for generating a SPARQL query, which can be made to OWN-PT for retrieving the possible answers.

3.3 Fixed Argument Extraction

Besides classifying the question into a relation type, the fixed relation argument x must be extracted from the input text. In all handcrafted templates, this argument is the last term of the question. In fact, for the type of considered questions, there would not be many variations where

this was not the case. Therefore, the extraction of x was simplified in such a way that it is always the last sequence of words in the question. More precisely, in order to cover multiword expressions, the system searches for the longest lexical form in OWN-PT starting with the i^{th} , $i \in (1, n]$, and ending in the last token of the question. For instance, given the question *que tipos há de intoxicação alimentar?* (what types are there of food poisoning?), the system checks, in the following order, whether OWN-PT covers the forms: *tipos há de intoxicação alimentar*, *há de intoxicação alimentar*, *de intoxicação alimentar*, *intoxicação alimentar*. It stops once it finds that the lexical form *intoxicação alimentar* (food poisoning) exists.

3.4 SPARQL Generation

With the label and the fixed argument, a SPARQL query can be generated to get all the valid lexical forms for y . Figure 1 shows the generated query for the question *que formas há de correr?*, with label [hyponymOf] and $x = correr$. It retrieves lexical forms (lf) in OWN-PT synsets ($s2$) for which the aligned PWN synset ($sen2$) is a hyponym of another PWN synset ($sen1$) that is aligned with an OWN-PT synset with the lexical form *correr*.

```
prefix wn30: <https://w3id.org/own-pt/wn30/schema/>
prefix owl: <http://www.w3.org/2002/07/owl#>

SELECT ?lf WHERE {
  ?spt1 wn30:containsWordSense ?ws1 .
  ?ws1 wn30:word ?word .
  ?word wn30:lexicalForm "correr"@pt .
  ?sen1 owl:sameAs ?spt1 .
  ?sen2 owl:sameAs ?spt2 .
  ?sen2 wn30:hyponymOf ?sen1 .
  ?spt2 wn30:containsWordSense ?ws2 .
  ?ws2 wn30:word/wn30:lexicalForm ?lf .
}
```

Figure 1: SPARQL query for retrieving the hyponyms of *correr*. Query is available in OWN-PT’s SPARQL endpoint at <https://ibm.co/20Cptyv>.

4 Experiments

NELIO was implemented in Java, using Apache Jena² for querying OWN-PT and DeepLearning4J³ for training the doc2vec model, more specifically, the ParagraphVectors class. This sections illustrates NELIO’s usage and reports on a simple evaluation made automatically.

²<https://jena.apache.org/>

³<https://deeplearning4j.org/>

X hyponymOf Y (8 templates)	
que (tipos géneros espécies sub-classes especificações formas) há de <Y>?	<i>what (types genres species subclasses specifications forms) are there of (Y)?</i>
que (hipónimos subordinados) tem <Y>?	<i>what (hyponyms subordinates) does (Y) have?</i>
X hypernymOf Y (4 templates)	
qual é a classe de <Y>?	<i>what is the class of (Y)?</i>
qual é o hiperónimo de <Y>?	<i>what is the hypernym of (Y)?</i>
qual é o conceito superordenado de <Y>?	<i>what is the superordinate concept of (Y)??</i>
o que é <Y>?	<i>what is (Y)?</i>
X memberHolonymOf Y (10 templates)	
quais os (membros constituintes componentes) de <Y>?	<i>what are the (members constituents components) of (Y)?</i>
que membros tem <Y>?	<i>what members does (Y) have?</i>
o que tem <Y>?	<i>what does (Y) have?</i>
de que é constituído <Y>?	<i>what is (Y) made of?</i>
o que inclui <Y>?	<i>what does (Y) include?</i>
o que está em <Y>?	<i>what is there in (Y)?</i>
em que se (divide decompõe) <Y>?	<i>in what can (Y) be divided decomposed?</i>
X partHolonymOf Y (9 templates)	
quais as (partes constituintes componentes) de <Y>?	<i>what are the (parts constituents components) of (Y)?</i>
que partes tem <Y>?	<i>what parts does (Y) have?</i>
o que tem <Y>?	<i>what does (Y) have?</i>
de que é constituído <Y>?	<i>what is (Y) made of?</i>
o que inclui <Y>?	<i>what does (Y) include?</i>
em que se (divide decompõe) <Y>?	<i>in what can (Y) be (divided decomposed)?</i>
X substanceHolonymOf Y (7 templates)	
quais as substâncias de <Y>?	<i>what are the substances of (Y)?</i>
que substâncias tem <Y>?	<i>what substances does (Y) have?</i>
o que tem <Y>?	<i>what does (Y) have?</i>
de que é (constituído feito) <Y>?	<i>what is (Y) made of?</i>
o que inclui <Y>?	<i>what does (Y) include?</i>
o que está em <Y>?	<i>what is there in (Y)?</i>
X memberMeronymOf Y / X partMeronymOf Y (3 templates)	
de que faz parte <Y>?	<i>what is part of (Y)?</i>
onde se inclui <Y>?	<i>where is (Y) included?</i>
a que pertence <Y>?	<i>what does (Y) belong to?</i>
X substanceMeronymOf Y (4 templates)	
de que faz parte <Y>?	<i>what is part of (Y)?</i>
onde se inclui <Y>?	<i>where is (Y) included?</i>
onde encontramos <Y>?	<i>where can we find (Y)?</i>
onde se encontra <Y>?	<i>where is (Y) found?</i>
X causes Y (7 templates)	
qual é o (efeito resultado) de <X>?	<i>what is the (effect result) of (X)?</i>
qual é a consequência de <X>?	<i>what is the consequence of (X)?</i>
o que (causa faz origina) <X>?	<i>what does (X) (cause make originate)?</i>
em que resulta <X>?	<i>what does (X) result in?</i>
X causes Y (2 templates)	
o que leva a <Y>?	<i>what leads to (Y)?</i>
o que resulta em <Y>?	<i>what does (Y) result in?</i>
X entails Y (4 templates)	
o que (acarreta implica) <Y>?	<i>what does (Y) (entail implies)?</i>
o que se (infere conclui) de <Y>?	<i>what may one (infer conclude) of (Y)?</i>
X antonymOf Y (5 templates)	
qual é o (antónimo contrário oposto inverso) de <X>?	<i>what is the (antonym contrary opposite inverse) of (Y)?</i>
o que é diferente de <X>?	<i>what is different from (Y)?</i>

Table 1: Handcrafted templates for querying the NLI interface. The vertical bar is an abbreviation, in the templates files, each word used constitute one template.

4.1 Examples

The resulting doc2vec model identifies the correct relation type in most situations. Besides being easy to train, a good thing about it is that no syntactic analysis is required and, still, the text of the questions does not have to match the original templates exactly. This provides an interesting level of tolerance while dealing with syntactic variability. Table 2 shows some of the questions answered correctly that, despite sharing some words, do not perfectly match any template.

On the other hand, since the model is based on the surface forms of words, it does not handle variation of vocabulary. For instance, the question “*qual é a matéria de chocolate?*” (what is the

matter of chocolate?) is answered with hyponyms of ‘chocolate’, instead of its substances, because no substanceHolonymOf template includes the word *matéria* (matter). However, this situation can be easily fixed by adding a single template with this word for the previous relation, e.g., *matéria-prima de <Y>?*.

4.2 Generalisation

Combining all possible question templates with all the relation instances in OWN-PT results in 298,249 different questions. Those questions were generated, shuffled and used as a dataset for training and testing NELIO.

In a simple experiment for analysing how well

Question	Answer
<i>tipos de peixe?</i> types of fish?	<i>sardinha, arenque, enchova, Peixe Esportivo, atum, salmoneiro, manjuba, eperlano, Salmão ...</i> sardine, herring, anchovy, sport fish, tuna, surmullet, smelt, salmon ...
<i>forma de andar?</i> way of walking?	<i>andar a galope, retornar, afastar-se, serpear, correr terras, descarrilhar, perseguir, coxear ...</i> gallop, return, depart, creep, drift, derail, seek, limp, ...
<i>o que é um cão?</i> what is a dog?	<i>animal doméstico, canídeo, cão</i> domestic animal, canid
<i>membros de Passeridae?</i> members of Passeridae?	<i>Pardal, pardal, Passer</i> sparrow, Passer
<i>que tem na Polónia?</i> what is there in Poland?	<i>Auschwitz-Birkenau, Prússia, Rio Vístula, Bydgoszcz, Rio Oder, Czeszochowa, Varsóvia, ...</i> Auschwitz-Birkenau, Prussia, Vístula river, Bydgoszcz, Oder river, Czeszochowa, Warsaw, ...
<i>onde fica a Polónia?</i> where is Poland?	<i>Europa</i> Europe
<i>que resulta de estender?</i> what results from stretching?	<i>esticar</i> to stretch
<i>que implica olhar?</i> what implies looking?	<i>olhar, ver, mirar, inspecionar, assistir, examinar, observar</i> to see, to eye, to inspect, to watch, to skim, to observe, to lay eyes on
<i>contrário de alto?</i> opposite of tall?	<i>baixo</i> short

Table 2: Questions correctly answered by NELIO.

the model generalises, it was tested with different proportions of training and testing data. Table 3 presents the accuracy, i.e., the proportion of questions correctly answered in this experiment.

This also showed that, some of the incorrect answers were in fact empty, due to misclassification of the relation type, which suggested a second experiment: similar to the previous but, when the given answer was empty, NELIO tried to get an answer with the second or third relation type predicted by doc2vec. As expected, this resulted in higher accuracies, also in table 3 (Top-3).

Training		Test Prop.	Accuracy	
Prop.	#Questions		1st label	Top-3
90%	(268,424)	10%	93.3%	97.2%
75%	(223,687)	25%	92.9%	97.5%
50%	(149,125)	50%	93.2%	97.6%
25%	(74,562)	75%	91.9%	97.6%
20%	(59,650)	80%	92.1%	97.9%
15%	(44,737)	85%	89.6%	97.0%
10%	(29,825)	90%	80.3%	95.9%
5%	(14,912)	95%	79.2%	94.9%

Table 3: Accuracy when answering questions depending on proportion of training data.

When considering only the top label, training the model with 90% ($\approx 268k$), 50% ($\approx 149k$), or even 20% ($\approx 59k$) of the questions, results in accuracies above 90%. This happens mainly because, although there are only a few templates, they are instantiated many times. With lower training proportions, accuracy drops more considerably. Yet, with only 5% it is still close to 80%.

Accuracy is different for different relations. For instance, with 90% of training data, it ranges from 100%, for entails, antonymOf and hyponymOf, to

73%, for substanceMeronymOf. A closer look shows that, except for the meronym-holonym relations, all accuracies are higher than 94% (hypernymOf). The problem with the former is that they are very similar and, for this reason, share several templates among them, which confuses the model.

The aforementioned issue is significantly minimised when the top-3 labels are considered. In this case, accuracies are 97% or higher with 15% or more training data. Specifically, they are 98% or higher for all relation types, except for the meronym-holonym, which are still the most problematic. The lower accuracy in this scenario is for memberHolonymOf (87.9%).

5 Related Work

Traditional Automatic Question Answering (QA) follows an Information Retrieval perspective (Kolomyets and Moens, 2011). Queries are typically natural language questions (NLQs) and answers are retrieved from a collection of written documents. But the development of natural language interfaces (NLIs) for databases has also been a research topic for a long time (Androutsopoulos et al., 1995). Here, the primary challenge involves translating NLQs to formal queries made to a database. Knowledge-based QA systems are a specific case of the previous.

Several NLIs for ontologies — e.g., Querix (Kaufmann et al., 2006), PANTO (Wang et al., 2007), FREyA (Damjanovic et al., 2010) — translate NLQs to SPARQL with a set of rules on the result of syntactically parsing NLQs, possibly using PWN for synonym expansion. A

similar approach (Unger et al., 2012) may be based on SPARQL templates, to be filled with entities and predicates identified in the NLQ.

Other systems rely on domain-independent semantic parsers that learn how to map NLQs to predicates in a large knowledge base, based on question-answer pairs. SEMPRE (Berant et al., 2013) maps words to predicates and then combines the predicates to the final logical form. Another possibility (Kwiatkowski et al., 2013) is to parse utterances for producing an underspecified logical form, before mapping lexical predicates to the target ontology predicates. The previous systems were assessed while resorting to Freebase for answering NLQs. Yet, as opposing to Freebase or DBpedia, wordnets have a much smaller number of predicates. So, it could be worth exploring how semantic parsers could be adapted for our work.

Once translated to SPARQL, generally to a subset of this language, expressiveness is limited. To avoid this, SQUALL (Ferré, 2014) is a controlled natural language for querying and updating RDF datasets. Nouns and intransitive verbs are used as classes; relation nouns and transitive verbs as properties; and proper nouns as resources. Syntactic and semantic analysis is implemented as a Montague grammar, an approach that would work for querying a wordnet, considering the simplicity of its RDF model. On the other hand, SQUALL requires that end-users comply with its controlled syntax, and know the RDF vocabulary.

An alternative approach (Bordes et al., 2014) learns low-dimensional embeddings of words and entities, respectively in questions and relation types of Freebase. This way, representations of questions and of their corresponding answers are close to each other in the joint embedding space. More recent works (Neelakantan et al., 2016; Zhong et al., 2017) rely on neural networks for translating NLQs to formal queries, thus avoiding domain-specific grammars or rules.

6 Conclusion

We have described how we can leverage on sentence embeddings in the development of a NLI for a wordnet. The proposed procedure was applied to OWN-PT with some success. When trained in a subset with at least 20% of the possible questions, generated with a small set of templates, and tested with the remaining questions, accuracies were higher than 91%, when using the first pre-

diction, or 97%, when trying with the first three predictions, in case the previous did not return an answer. This simple experiment confirmed that the proposed approach works well with the doc2vec model for predicting the correct relation type. Despite the positive results, this experiment revealed that the system is confused by similar relations, for which the templates share vocabulary, namely the three types of meronymy. The problem can be minimised by considering the top-3 predictions, but others, such as merging the three relations, can be analysed in the future.

Still, this was a limited experiment, where known limitations of the system had a low impact. This includes questions with vocabulary not covered by the templates, or questions that do not end with the fixed word. The former can be minimised by adding alternative templates. The second is due to a simplification that works for many cases, but fails for some, as in the question ‘*quais frutas existem?*’ (what fruits exist?), where the target word is *frutas*. The previous question has to be made like ‘*quais os tipos de fruta?*’. In the future, we will devise more general ways of extracting the target argument from the question, e.g., having in mind that, among the words/expressions in the question, it should be the least frequent in the dataset; or maybe training an automatic sequence labeller for identifying the target argument in the context of a question. In the latter case, training data should also include templates that do not end with the target argument.

Other possible directions for future work include: (i) Presenting the answers according to the senses they apply to, because context is not enough for disambiguation (currently, there is an option for considering only the first sense); (ii) Adding alternative types of question e.g., *what is the relation between $\langle x \rangle$ and $\langle y \rangle$? or is $\langle y \rangle$ related to $\langle x \rangle$?*, to be answered, respectively, with the name of a relation between x and y in OWN-PT, or yes/no, depending on the existence of such a relation; (iii) Exploring recent models for representing sentence meaning, learned from natural language inference data (Conneau et al., 2017), though available data in Portuguese (Fonseca et al., 2016; Real et al., 2018) may not be enough.

Despite its limitations, NELIO was already helpful for finding issues in OWN-PT that need to be fixed. It showed flaws such as inconsistencies

in the capitalization (e.g., *Salmão, Pardal*), presence of underscores instead of spaces (e.g., *animal doméstico*), or plural instead of singular form (e.g., *epidemias, montanhas*), not to mention actual errors (e.g., *dançar* entails *andar*, in English, dancing entails walking).

A mid-term goal is to make NELIO available from a web interface. In the meantime, its source code is available online, at https://github.com/hgoliv/nli_openwordnet-pt. Although, so far, the proposed approach was only used as a NLI for a wordnet, in principle, a similar approach could be used in the development of a NLI for any knowledge base represented as *a*-related-to-*b* triples.

Acknowledgments

This work was partially funded by FCT's INCoDe 2030 initiative, in the scope of the demonstration project AIA, "Apoio Inteligente a empreendedores (chatbots)".

References

- I. Androutsopoulos, G.D. Ritchie, and P. Thanisch. 1995. [Natural language interfaces to databases — an introduction](#). *Natural Language Engineering*, 1(1):29–81.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544. ACL Press.
- Francis Bond and Ryan Foster. 2013. Linking and extending an open multilingual wordnet. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1352–1362. ACL Press.
- Antoine Bordes, Sumit Chopra, and Jason Weston. 2014. [Question answering with subgraph embeddings](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 615–620. ACL Press.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. ACL Press.
- Danica Damjanovic, Milan Agatonovic, and Hamish Cunningham. 2010. [Natural language interfaces to ontologies: Combining syntactic analysis and ontology-based lookup through the user interaction](#). In *Proceedings of the 7th International Conference on The Semantic Web: Research and Applications - Volume Part I, ESWC'10*, pages 106–120. Springer.
- Christiane Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database (Language, Speech, and Communication)*. The MIT Press.
- Sébastien Ferré. 2014. [Squall: a controlled natural language as expressive as SPARQL 1.1](#). *Data Knowl. Eng.*, 94(PB):163–188.
- Erick Rocha Fonseca, Leandro Borges dos Santos, Marcelo Criscuolo, and Sandra Maria Aluísio. 2016. Visão geral da avaliação de similaridade semântica e inferência textual. *Linguamática*, 8(2):3–13.
- Esther Kaufmann, Abraham Bernstein, and Renato Zumstein. 2006. Querix: A natural language interface to query ontologies based on clarification dialogs. In *Proceedings of the 5th International Semantic Web Conference (ISWC 2006)*, pages 980–981.
- Oleksandr Kolomiyets and Marie-Francine Moens. 2011. A Survey on Question Answering Technology from an Information Retrieval Perspective. *Information Sciences*, 181(24):5412–5434.
- Tom Kwiatkowski, Eunsol Choi, Yoav Artzi, and Luke Zettlemoyer. 2013. [Scaling semantic parsers with on-the-fly ontology matching](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1545–1556, Seattle, Washington, USA. ACL Press.
- Quoc Le and Tomas Mikolov. 2014. [Distributed representations of sentences and documents](#). In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML'14*, pages II–1188–II–1196. JMLR.org.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *Proc. of the Workshop track of the Intl. Conf. on Learning Representations (ICLR)*, Scottsdale, Arizona.
- Arvind Neelakantan, Quoc V. Le, Martín Abadi, Andrew McCallum, and Dario Amodei. 2016. [Learning a natural language interface with neural programmer](#). *CoRR*, abs/1611.08945.
- Valeria de Paiva, Fabricio Chalub, Livy Real, and Alexandre Rademaker. 2016a. Making virtue of necessity: a verb lexicon. In *Proceedings of the 12th International Conference on the Computational Processing of Portuguese (PROPOR 2016)*, pages 271–282, Tomar, Portugal. Springer.
- Valeria de Paiva, Cláudia Freitas, Livy Real, and Alexandre Rademaker. 2014. [Improving the verb lexicon of OpenWordnet-PT](#). In *Proceedings of Workshop on Tools and Resources for Automatically Processing Portuguese and Spanish (ToRPorEsp)*, São Carlos, Brazil. Biblioteca Digital Brasileira de Computação, UFMG, Brazil.

- Valeria de Paiva, Alexandre Rademaker, and Gerard de Melo. 2012. [OpenWordNet-PT: An open Brazilian Wordnet for reasoning](#). In *Proceedings of COLING 2012: Demonstration Papers*, pages 353–360, Mumbai, India. The COLING 2012 Organizing Committee.
- Valeria de Paiva, Livy Real, Hugo Gonçalo Oliveira, Alexandre Rademaker, Cláudia Freitas, and Alberto Simões. 2016b. An overview of portuguese wordnets. In *Global Wordnet Conference 2016*, Bucharest, Romania.
- Alexandre Rademaker, Valeria de Paiva, Gerard de Melo, Livy Real, and Maira Gatti. 2014. [OpenWordNet-PT: A project report](#). In *Proceedings of the 7th Global WordNet Conference*, Tartu, Estonia.
- Livy Real, Fabricio Chalub, Valeria de Paiva, Claudia Freitas, and Alexandre Rademaker. 2015. [Seeing is correcting: curating lexical resources using social interfaces](#). In *Proceedings of the 4th Workshop on Linked Data in Linguistics: Resources and Applications*, pages 20–29, Beijing, China. ACL Press.
- Livy Real, Valeria de Paiva, Fabricio Chalub, and Alexandre Rademaker. 2016. Gentle with gentilics. In *Joint Second Workshop on Language and Ontologies (LangOnto2) and Terminology and Knowledge Structures (TermiKS) (co-located with LREC 2016)*, Slovenia.
- Livy Real, Ana Rodrigues, Andressa Vieira, Beatriz Albiero, Bruna Thalenberg, Bruno Guide, Cindy Silva, Guilherme de Oliveira Lima, Igor C. S. Câmara, Miloš Stanojević, Rodrigo Souza, and Valeria De Paiva. 2018. Sick-br: A portuguese corpus for inference. In *13th International Conference, PROPOR 2018*, pages 303–312, Canela, Brazil.
- Christina Unger, Lorenz Bühmann, Jens Lehmann, Axel-Cyrille Ngonga Ngomo, Daniel Gerber, and Philipp Cimiano. 2012. [Template-based question answering over rdf data](#). In *Proceedings of the 21st International Conference on World Wide Web, WWW '12*, pages 639–648, New York, NY, USA. ACM.
- Chong Wang, Miao Xiong, Qi Zhou, and Yong Yu. 2007. [Panto: A portable natural language interface to ontologies](#). In *Proceedings of the 4th European Conference on The Semantic Web: Research and Applications, ESWC '07*, pages 473–487. Springer.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*.