



HAL
open science

Automatic Extension of WOLF

Benoît Sagot, Darja Fišer

► **To cite this version:**

Benoît Sagot, Darja Fišer. Automatic Extension of WOLF. GWC2012 - 6th International Global Wordnet Conference, Global Wordnet Association + Toyohashi University of Technology + National Institute of Japanese Language and Linguistics, Jan 2012, Matsue, Japan. hal-00655774

HAL Id: hal-00655774

<https://inria.hal.science/hal-00655774>

Submitted on 2 Jan 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Automatic Extension of WOLF

Benoît Sagot

INRIA / University Paris 7
Paris, France
benoit.sagot@inria.fr

Darja Fišer

University of Ljubljana
Ljubljana, Slovenia
darja.fiser@ff.uni-lj.si

Abstract

In this paper we present the extension of WOLF, a freely available, automatically created wordnet for French, the biggest drawback of which has until now been the lack of general concepts that are typically expressed with highly polysemous vocabulary that is on the one hand the most valuable for applications in human language technologies but also the most difficult to add to wordnet accurately with automatic methods on the other. Using a set of features, we train a Maximum Entropy classifier on the existing core wordnet to be able to assign appropriate synset ids to new words, extracted from multiple, multilingual sources of lexical knowledge, such as Wiktionaries, Wikipedias and corpora. Automatic and manual evaluation shows high coverage as well as high quality of the resulting lexico-semantic repository of. Another important advantage of the approach is that it is fully automatic and language-independent and could therefore be applied to any other language still lacking a wordnet.

1 Introduction

Whatever the framework and representation of lexical knowledge, such as ontologies, framenets or wordnets, semantic lexicons can only contribute to applications in human language technologies, such as word-sense disambiguation, information extraction or machine translation, if their coverage is comprehensive as well as accurate. The language resources development community seems to have reached a consensus that, despite giving the most reliable results, manual construction of lexical resources is to time-consuming and expensive to be practical for most purposes. Several semi- or fully automatic approaches have been proposed instead, exploiting various types of existing resources to facilitate the development of a new semantic lexicon, especially

wordnets. However, most proposed approaches to induce a wordnet automatically, still suffer from the necessary trade-off between limited coverage and the desired level of accuracy, both of which are required if the resource is to be useful in a practical application.

This is why we present here an approach for wordnet extension by extracting additional lexico-semantic information from already available bilingual language resources and then training a maximum entropy classifier on the existing core wordnet in order to assign the new vocabulary to the appropriate synsets. Our approach, applied on the French wordnet WOLF, is comprehensive in that it can handle monosemous and polysemous words from all parts of speech which belong to the general vocabulary as well as specialized domains and can also deal with multi-word expressions.

The rest of the paper is structured as follows: in Section 2 we give an overview of related work. In Section 3 we introduce the current edition of WOLF. In Section 4, we describe the process of extracting lexico-semantic information from bilingual lexical resources. In Section 5 we explain the wordnet enrichment experiment using a maximum entropy classifier that helped us determine whether a translation we extracted from the existing resources is an appropriate candidate for a given synset. Section 6 is dedicated to the analysis and evaluation of the extended resources, and Section 7 contains concluding remarks and ideas for future work.

2 Related work

Most automatic approaches to create a wordnet for a new language take the Princeton wordnet as a backbone and extend it with the vocabulary inventory of the target language. One of the most straightforward and most widely used resources to obtain lexical knowledge for the language in question are machine-readable bilingual diction-

aries. Entries from the dictionary are linked to PWN synsets under the assumption that their counterparts in the target language correspond to the same synset (Knight and Luk 1994). A well-known problem with this approach is that bilingual dictionaries are generally not concept-based but follow traditional lexicographic principles, which is why the biggest obstacle is the disambiguation of dictionary entries.

When such dictionaries are not available or when they do not contain sufficient information to disambiguate the entries, bilingual lexicons can be extracted from parallel corpora (Fung 1995). The underlying assumption here is that senses of ambiguous words in one language are often translated into distinct words in another language (Dyvik 2002). Furthermore, if two or more words are translated into the same word in another language, then they often share some element of meaning (Ide et al. 2002). This results in sense distinctions of a polysemous source word or yields synonym sets.

The third set of approaches that have become popular in the past few years extract the meaning, translations and relationships between words in one or several languages from Wikipedia. New wordnets have been induced by using structural information to assign Wikipedia categories to WordNet (Ponzetto and Navigli 2009) or by extracting keywords from Wikipedia articles (Reiter et al. 2008). Vector-space models to map Wikipedia pages to Wordnet have been developed (Ruiz-Casado et al. (2005)). The most advanced approaches use Wikipedia and related projects, such as Wiktionary, to bootstrap wordnets for multiple languages (Melo and Weikum 2009, Navigli and Ponzetto 2010).

An unsupervised machine-learning approach has been used by Montazery and Faili (2011) to construct a wordnet for Persian. Their approach is similar to ours in the sense that they too combine translation candidates obtained from bilingual dictionaries and corpus-based contextual information for these candidates to establish links to Princeton WordNet synsets.

3 Wordnet Libre du Français (WOLF)

Previous work on the development of WOLF (Fišer and Sagot 2008) has focused on benefiting from available resources of three different types: general and domain-specific bilingual dictionaries, parallel corpora and Wiki resources (Wikipedia and Wiktionaries).

The core WOLF was created by disambiguating (*literal, synset*) pairs obtained from a word-aligned multilingual parallel corpus with the help of already existing wordnets for several languages other than French. For each multilingual lexicon entry, translation equivalents in all these languages were assigned a set of possible synset ids from their wordnets. Assuming that translation equivalents in the word-aligned parallel corpus are lexicalizations of the same concept, they shared one or several intersecting synset ids which were then also assigned to the French equivalent in the lexicon. The approach was limited to almost predominantly basic synsets which were common among all the wordnets used, and to single-word literals because cross-lingual mapping of multi-word expressions was not possible with or word-alignment procedure. In order to compensate these shortcomings, additional (*literal, synset*) pairs for monosemous English words were also harvested from various freely-available bilingual resources (dictionaries, thesauri and Wikipedia). Sense assignment for these was near perfect because they did not require any disambiguation.

The wordnet for French created in this way contained about 32,300 non-empty synsets, 87% of which were nominal. With the approach we adopted, we were able to populate just over 50% of Base Concept Sets and 25% of the rest of the concepts from Princeton WordNet. The first version of WOLF was already bigger than the French WordNet (22,121 synsets) that had been developed within the EuroWordNet Project (Vossen 1999) and was comparable to the more recent wordnet construction contribution called JAWS (34,367 synsets) which was developed by Mouton and de Chalendar (2010) from a bilingual dictionary, which however contains only nouns.

Manual evaluation of the results showed that the wordnet generated in this way is relatively reliable but does not use full potential of the available resources. This is why we have devised an additional large-scale extension cycle, aiming at taking full advantage of the existing lexical resources in order to improve the coverage of WOLF without compromising its accuracy while the first version of WOLF will serve as the baseline. The procedure is described in the rest of this paper. We begin by presenting the various resources used in the experiment and the way we extracted (*literal, synset*) candidates from them. Then we introduce the maximum entropy classifier and the features we use for filtering these

pairs and extending our initial wordnet. We also report manual and automatic evaluation of the results and look into possible steps to refine the developed resource in the future.

4 Bilingual lexicon extraction

In this experiment we used two types of sources of lexical knowledge: the structured freely-available general and domain specific bilingual dictionaries, and the semi-structured articles from the on-line Wikipedia. The main goal of the extraction process was to extract as many French translation variants for each English word as possible in order to capture as many senses of that word as possible. With this we obtained *wordnet candidates* in the form of (*literal, synset*) pairs, i.e. a French translation of an English word with an assigned synset id from Princeton WordNet.

General vocabulary was extracted from the English and French **Wiktionary** in which translations are explicitly encoded for all parts-of-speech. The number of pairs extracted from each resource is given in Table 1. For domain-specific vocabulary we used **Wikispecies**, a taxonomy of living species that includes both Latin standard names and vernacular terms.

Less structured than dictionaries but with a much more predefined structure than free text is the on-line multilingual collaborative encyclopaedia **Wikipedia**. We used English and French articles by following inter-language links that relate two articles on the same topic. We enhanced the extraction process with a simple analysis of article bodies with which we resolved ambiguities arising from the capitalization of article titles (e.g. *Grass-novelist, Grass-plant*). In a similar way we also identified synonyms for the article titles (e.g. “*Cannabis*, also known as *marijuana*”), their definitions (e.g. “*Hockey* is a family of sports in which two teams play against each other by trying to manoeuvre a ball or a puck into the opponent’s goal using a hockey stick.”) and usage examples.

Resources used	En-Fr equivalents
English Wiktionary	39,286
French Wiktionary	59,659
Wikispecies	48,046
Wikipedia	286,818
Total (duplicates removed)	417,419

Table 1: Results of bilingual lexicon extraction from heterogeneous resources

The result of our extraction process is a large bilingual lexicon of all English-French translation pairs with the name of the resource they originate from. The figures for both extracted bilingual lexicons are summarized in Table 1.

As can be seen in Table 1, we were able to extract a substantial amount of bilingual entries from the various resources. However, the extracted entries suffer from an important drawback: they do not contain any explicit information that can help us map these entries to PWN, neither do they contain contextual information from corpus occurrences that would help us determine their sense based on their usage. For example, an English-French translation pair (*dog, chien*), which we extracted from the Wiktionary, does not contain any information that would make it possible for us to determine which of the 8 synsets in WPN containing the literal *dog* would be appropriate to be translated with *chien* in WOLF. In Wiktionary articles, translations of a given word are sometimes organized into senses and described with short glosses. These have been compared to PWN glosses in order to map Wiktionary senses to PWN synsets (see Bernhard and Gurevych 2009). The first sentence of a Wikipedia article can be used in a similar way (see Ruiz-Casado et al. 2005). However, this is not the case for all Wiktionary entries or for other resources. Therefore, at this point, we assign to each translation pair *all* possible synset ids and disambiguate it in the next step.

5 Large-scale wordnet extension

Restricting the use of a bilingual lexicon to monosemous English literals is a safe but very limited approach that does not exploit the available resources to their full potential, which is a waste of resources and should be improved. However, using lexicon-based candidates generated from polysemous English literals is only possible if we can establish the likelihood with which a word should be added to a particular synset, i.e. can compute the semantic distance between a given French literal and PWN synset id.

In this paper we propose a technique to achieve exactly that. It is based on the core version of WOLF and a probabilistic classifier that uses various features associated with each (*literal, synset*) candidate.

5.1 Training set and feature selection

First we extract all (*literal, synset*) pairs from the bilingual lexicon that are already in the baseline wordnet and consider them as valid ones (score 1). All the other candidates, on the other hand, are considered invalid (score 0). This creates a noisy but reasonable training set for a probabilistic model. It is noisy for two reasons: first, our baseline wordnet does contain some mistakes because synsets were generated automatically and have not been completely manually validated; second, and more important reason is the fact that the baseline wordnet is not complete, which is why our new candidates may be valid even though they are not present in the baseline wordnets. It is precisely these candidates we are looking for in our wordnet extension procedure.

In order to use the baseline wordnet as a training set for our classifier that will assign scores to all the candidates in the lexicon, we need to extract features from (*literal, synset*) pairs in WOLF.

5.1.1 Semantic proximity

The central feature we use models the **semantic proximity** between a literal and a synset. The feature can be illustrated on the example (*dog, chien*) we already used above. There are 8 PWN synsets that contain the literal *dog*, which is why this bilingual entry yields 8 different (*literal, synset*) candidates. We now need to determine which of these 8 candidates are valid. In other words, we need to establish which of the 8 corresponding synsets the French literal *chien* should be added to in WOLF. We therefore compute the semantic similarity of the literal *chien* w.r.t. each of these 8 synsets. For doing this, we first represent each WOLF synset by a bag of words obtained by extracting all literals from this synset and all the synsets up to 2 nodes apart in WOLF. For example, the synset {*andiron, firedog, dog, dog-iron*} in PWN, which is empty in the baseline WOLF, is represented by the bag of words {*appareil, mécanisme, barre, rayon, support, balustre, ...*} (~*device, mechanism, bar, shelve, baluster, ...*). Next, we use a distributional semantic model for evaluating the semantic similarity of *chien* w.r.t. this bag of words. We use the freely-available SemanticVectors package (Widdows and Ferraro 2008). The distributional semantic model was built from the 65,000 lemmatised webpages from the French web corpus frWaC corpus (Ferraresi *et al.* 2010). This gives us a semantic similarity score between *chien* and the synset {*andiron, firedog, dog, dog-iron*},

which is only 0.035, while the similarity between *chien* and one of its valid synsets, {*dog, domestic dog, Canis familiaris*} is as high as 0.331.

5.1.2 Additional features

In addition to semantic proximity, we use a number of other supporting features which are described below. Let us consider a candidate (*T, S*) that has been generated because our bilingual resources provided entries of the form (E_1, T)...(E_n, T), where all PWN literals E_i 's are among *S*'s literals. **The number of such PWN literals** is one of the features. **Each possible source** (e.g. English Wiktionary) corresponds to one feature, which receives the value 1 if and only if at least one of the (E_i, T) entries was extracted from this source. We also extract **the lowest polysemy index** among all E_i 's: if one of the E_i 's is monosemous, this feature receives the value 1; if the least polysemous E_i is in two PWN synsets, this feature receives the value 2. The idea is that if the candidate is generated from at least one monosemous PWN literal, it is very likely to be correct, whereas if it was generated from only highly polysemous PWN literals, it is much more questionable. Finally, **the number of tokens** in *T* is used as a feature (often, literals with many tokens are not translations of PWN literals but rather glosses).

5.2 Classifier training

Based on these features, we train a classifier using the Maximum-Entropy package *megam* (Hal Daumé III, 2004). An analysis of the models shows that the semantic similarity is by far the strongest feature. As expected, the lowest polysemy index among English literals also contributes positively, as does the number of different English literals yielding the generation of the candidate, and the number of sources involved. On the other hand, also as expected, the number of tokens in the target language literal has a negative impact on the certainty score.

The result of our classifier on a given (*literal, synset*) candidate is a score between 0 (bad candidate) and 1 (good candidate). We empirically set the threshold at 0.1 (see Section 6.1) for adding the candidate to the wordnet. The results are presented and evaluated in the next section.

6 Results and evaluation

6.1 Analysis of the results

Our wordnet extension procedure yielded 55,159 French wordnet candidates (out of 177,980).

Among the 55,159 French candidates, 15,313 (28%) correspond to (*literal, synset*) pairs already present in the previous version of WOLF, which means that 39,823 (72%) new pairs were added. As a consequence, 13,899 synsets that were empty in the previous version of WOLF now have at least one French literal.

A comparison of WOLF before and after the extension paralleled with the figures from Princeton WordNet 3.0 is given in in Table 2. The extended version of WOLF has 43% more non-empty synsets than before the extension. The increase in the number of (*literal, synset*) pairs in the new WOLF is even higher; the number rose from 46,411 to 76,436 (+65%).

	PWN 3.0	WOLF old	WOLF new
N	82,114	28,559	36,933
V	13,767	1,554	4,105
Adj	18,156	1,562	4,282
Adv	3,621	871	1,125
Total	117,658	32,550	46,449
BCS1-3	4,671	4,339	6,171
Non-BCS	112,987	28,211	40,278

Table 2: Results of the wordnet extension procedure

As in PWN, by far the most frequent domain is Factotum, and the order for the following three most frequent domains is the same in both wordnets as well (Zoology, Botany, Biology). Most synsets belonging to these domains were generated from Wikispecies and Wikipedia while Wiktionary was the most frequent source for the Factotum domain. Of all the wordnet domains, only 3 are missing in WOLF (Paleontology, Rugby, and Volleyball) but these domains have less than 10 synsets in total even in PWN.

Average synset length in the extended WOLF is 1.79 literals per synset, which is slightly more than in PWN 3.0 (1.76). It is the lowest for nominal synsets (1.72) and the highest for adverbial ones (2.06). In PWN adverbial synsets are by far the shortest (1.54) while verbal ones are the longest (1.82). The longest synset in the extended WOLF is an adverbial one which contains as many as 27 literals, while in PWN the longest synset is a nominal one with 28 literals.

Table 3 contains a comparison between the level of polysemy when taking into account all literals vs. considering only polysemous ones. The comparison shows that while English literals are on average more polysemous than the French ones, there are big differences between English and French verbs, suggesting that automatically gen-

erated French verbal synsets contain some noise which will have to be filtered out in the future.

	PWN 3.0	WOLF new
avg. poly. + mono.	1.39	1.28
N	1.23	1.19
V	2.17	3.36
avg. poly. - mono	2.91	2.11
N	2.77	1.84
V	3.57	5.0

Table 3: Results of the wordnet extension procedure

A comparison of unique literals in PWN and WOLF shows that we were able to automatically generate as much as 25% of all multi-word expressions and over 30% of proper names found in PWN, which is a very good result, considering that the only source of both of these groups of literals was Wikipedia.

6.2 Manual evaluation of the results

In this section we report the results of manual evaluation of the wordnet extension where we evaluate the accuracy of the (*literal, synset*) candidates we obtained with the classifier as well as the accuracy of the candidates we discarded. For the evaluation we randomly selected 400 hundred (*literal, synset*) and evaluated them manually, using only two tags: “OK” if it would be correct to add that literal to the synset, and “NO” if it would be wrong, regardless of what the reason was for the error and how semantically close it was to the synset. The accuracy of a set of candidates is as usual as the proportion of candidates receiving the “OK” tag. Moreover, in order to assess the quality of our scoring technique, we compared the accuracy of the candidates per quartile w.r.t. their certainty scores.

The results of manual evaluation are shown in Table 4. They show a strong correlation between the certainty score they received and the accuracy of the candidates, thus justifying our decision to use this threshold but other threshold values could have been used too: higher values would have provided candidates with an even higher accuracy but the scale of the wordnet extension would have been lower; on the other hand, lower threshold values would have extended our wordnets even more, but would have introduced much more noise.

No. of candidates evaluated	400
No. of candidates added to wordnet	27%
Accuracy of all candidates	52%
Acc. of the candidates added to WOLF	81%

Accuracy of the discarded candidates	40%
Accuracy in the upper (4 th) quartile	83%
Accuracy in the third quartile	63%
Accuracy in the second quartile	41%
Accuracy in the lower (1 st) quartile	20%

Table 4: Manual evaluation of (*literal, synset*) candidates generated for extending WOLF

6.3 Automatic evaluation of the results

In this section we report the results of automatic evaluation of the generated wordnet against the already existing wordnet for French that was developed within the EuroWordNet project. With this evaluation we will gain an insight into the precision and recall of the wordnet we created with the proposed extension procedure. However, such an evaluation is only partial, because the detected discrepancies between the two resources are not only errors in our automatically created wordnets but can also stem from a missing literal in the resource we use for comparison. Automatic evaluation was performed on non-empty synsets, which means that adjectival and adverbial synsets in WOLF could not be evaluated this way at all because other existing French wordnets do not cover them.

When considering non-empty synsets in FWN, any (*literal, synset*) pair that is common to both resources is considered correct. When the number of valid (*literal, synset*) pairs of all types are combined, we reach a total of ~65,690 valid pairs out of 76,436, reaching a ~86% accuracy. A direct comparison to other related resources developed by Navigli and Ponzetto (2010) and di Melo and Weikum (2010) is not straightforward because even though the resources we used overlap to a great extent, their aim was to create a multilingual network while we focused only on French. An important difference between our approach and the one proposed by Navigli and Ponzetto (2010) is that they machine-translated the missing translations, while we only use resources that were created by humans, which is why we have more accurate translations. On the other hand, while di Melo and Weikum’s (2010) wordnet for French has a slightly higher accuracy, it is smaller than ours. This shows that the approach we used to benefit as much as possible from available resources using basic NLP tools only is very efficient for building large-scale reliable wordnets.

	Correct (<i>literal, synset</i>) pairs in WOLF and FWN	Correct WOLF pairs not in FWN	Incorrect WOLF pairs not in FWN	Correct FWN pairs not WOLF
Nominal pairs not empty in FWN	8,474	11,627		15,474
	correct pairs: ~15,915		incorrect pairs: ~4,186	
Verbal pairs not empty in FWN	1,826	3,859		6,168
	correct pairs: ~3,177		incorrect pairs: ~2,508	
Empty pairs in FWN	0	50,650		0
	correct pairs: ~46,598		incorrect pairs: ~4,052	
All pairs	10,300	66,136		21,642 + the no. of literals missing in synsets not covered by FWN
	~65,690 overall precision: ~86 %		incorrect pairs: ~10,746	

Table 5: Automatic evaluation of the extended WOLF based on FWN

7 Results and evaluation

In this paper we described an approach to extend an existing wordnet from heterogeneous re-

sources. Using various features such as distributional similarity, we were able to reuse automatically extracted bilingual lexicons for translating and disambiguating polysemous literals, which

had so far been dealt only with word-aligned corpora. The result of our work is a freely available lexical semantic resource that is large and accurate enough for use in real HLT applications. Compared to other similar resources for French, our wordnet is bigger than the much older French WuroWordNet and more comprehensive than the much more recent JAWS database. Due to the multiple human-produced resources which it was based on is more accurate than BabelNet (Navigli and Ponzetto, 2010) and larger than the French part of the multilingual wordnet developed by di Melo and Weikum (2010).

Analysis and evaluation of the approach shows that it is both versatile and accurate enough to successfully extend a wordnet of limited coverage. Another major advantage of the approach is that it is fully modular, adaptable and language independent and can therefore be used for any language still lacking a substantial wordnet.

In the future we plans to adapt the distributional similarity measure in order to automatically detect literas that are outliers in synsets and should therefore be removed from the developed wordnet. This procedure will provide an even more accurate and useful source of the much needed lexical knowledge that is much needed in virtually all HLT tasks.

Acknowledgments

The work described in this paper has been funded in part by the French-Slovene PHC PROTEUS project “Building Slovene-French linguistic resources: parallel corpus and wordnet” (22718UC), by the French national grant EDyLex (ANR-09-CORD-008) and by the Slovene national postdoctoral grant (Z6-3668).

References

Delphine Bernhard and Iryna Gurevych (2009). Combining Lexical Semantic Resources with Question & Answer Archives for Translation-Based Answer Finding. *Proc. ACL-IJCNLP'09*.

Hal Daumé III (2004). *Notes on CG and LM-BFGS optimization of logistic regression*.

Mona Diab (2004). The Feasibility of Bootstrapping an Arabic WordNet leveraging Parallel Corpora and an English WordNet. In *Proc. of NEMLAR*, Cairo, Egypt.

Helge Dyvik (2002). Translations as semantic mirrors: from parallel corpus to wordnet. Revised version of paper presented at *ICAME 2002*, Gothenburg, Sweden.

Christiane Fellbaum, editor. 1998. *WordNet: An Electronic Database*. MIT Press, Cambridge, Massachusetts, USA.

Adriano Ferraresi, Silvia Bernardini, Giovanni Picci and Marco Baroni (2010). Web Corpora for Bilingual Lexicography: A Pilot Study of English/French Collocation Extraction and Translation. In Xiao, R. (ed.) *Using Corpora in Contrastive and Translation Studies*. Newcastle: Cambridge Scholars Publishing.

Darja Fišer and Sagot Benoît (2008). Combining multiple resources to build reliable wordnets. In *Proc. of TSD'08*, Brno, Czech Republic.

Pascale Fung (1995). A pattern matching method for finding noun and proper noun translations from noisy parallel corpora. In *Proc. of ACL'95*.

Nancy Ide, Tomaz Erjavec, Dan Tufis (2002). Sense Discrimination with Parallel Corpora. In *Proc. of the ACL'02 Workshop on Word Sense Disambiguation*, Philadelphia, USA.

Kevin Knight and Steve K. Luk (1994). Building a large-scale knowledge base for machine translation. In *Proc. of AAAI'94*.

Gerard de Melo and Gerhard Weikum (2009). Towards a universal wordnet by learning from combined evidence. In *Proc. of CIKM'09*.

Mortaza Montazery and Hesham Faili (2011). Unsupervised Learning for Persian WordNet Construction. In: *Proc. of RANLP'11*, Hissar, Bulgaria.

Claire Mouton and Gaël de Chalendar (2010). JAWS: Just Another WordNet Subset. In *Proc. of TALN'10*, Montreal, Canada.

Roberto Navigli, Simone Paolo Ponzetto: BabelNet: Building a Very Large Multilingual Semantic Network. In *Proc. of ACL'10*, Uppsala, Sweden.

Simone Paolo Ponzetto and Roberto Navigli. 2009. Large-scale taxonomy mapping for restructuring and integrating Wikipedia. In *Proc. of IJCAI'09*.

Maria Ruiz-Casado, Enrique Alfonseca, Pablo Castells (2005). Automatic Assignment of Wikipedia Encyclopedic Entries to WordNet Synsets. In *Proc. of AWIC'2005*. pp. 380–386

Dan Tufis (2000): BalkaNet — Design and Development of a Multilingual Balkan WordNet. In *Romanian Journal of Information Science and Technology Special Issue*, 7(1-2).

Piek Vossen, editor. 1998. *EuroWordNet: A Multilingual Database with Lexical Semantic Networks*. Kluwer, Dordrecht, The Netherlands.

Dominic Widdows and Kathleen Ferraro (2008). Semantic vectors: a scalable open source package and online technology management application. In *Proc. of LREC'08*, Marrakech, Morocco.