A Survey on Automatically-Constructed WordNets and their Evaluation: Lexical and Word Embedding-based Approaches

Steven Neale
School of English, Communication and Philosophy
Cardiff University, Wales, United Kingdom
NealeS2@cardiff.ac.uk

Abstract

WordNets – lexical databases in which groups of synonyms are arranged according to the semantic relationships between them – are crucial resources in semantically-focused natural language processing tasks, but are extremely costly and labour intensive to produce. In languages besides English, this has led to growing interest in constructing and extending WordNets automatically, as an alternative to producing them from scratch. This paper describes various approaches to constructing WordNets automatically – by leveraging traditional lexical resources and newer trends such as word embeddings – and also offers a discussion of the issues affecting the evaluation of automatically constructed WordNets.

Keywords: WordNet, evaluation, automatic synset extraction, bilingual dictionaries, word embeddings

1. Introduction

WordNets (Fellbaum, 1998) are lexical databases in which open-class words – nouns, verbs, adjectives and adverbs – are stored as sets of synonyms or 'synsets' and linked to each other by various semantic relationships. These relationships include antonymy (opposites, i.e. 'wet' and 'dry'), hypernymy and hyponymy (types and types-of, i.e. 'animal' and 'cat'), and meronymy (parts-of, i.e. 'toe' as a meronym of 'foot'), among others. The *Princeton WordNet (PWN)*[1] is the original and pioneering English language WordNet – now totalling over 117,000 synsets in version 3.0 – and its format has become the gold standard for lexical databases representing meanings and concepts.

Besides its adoption as a standard for lexical databases, the PWN has become a vital resource across a range of semantic processing tasks – it has been used to produce gold-standard semantically-annotated corpora such as *SemCor*[2] and to determine the correct meanings of words in natural language processing tasks such as word-sense disambiguation, text summarization, and semantic textual similarity. Naturally, its extensive usage in English has inspired the construction of new WordNets in many other languages. The *Global WordNet Association (GWA)*[3], for example, was set up to provide a platform for discussing, sharing and connecting WordNets in any language, while large-scale research projects such as *EuroWordNet* (Vossen, 2004) and *MultiWordNet* (Pianta et al., 2002) have focused on aligning synsets from multiple WordNets in different languages.

However, while the PWN has been constructed, extended and refined over decades, building new WordNets from scratch is an enormous undertaking. Manual construction of course ensures accurate synsets covering as many concepts as possible, but it requires lexicographers to spend hours crafting synsets and is thus far too labour-intensive to be feasible in most cases. Inevitably, this has led to a growing and improving body of research into techniques and approaches for constructing and extending WordNets automatically.

In this paper, a number of approaches to automatically construct and extend WordNets are presented, taking into account both lexical and word embedding-based approaches. In addition, a discussion of the evaluation of these resources – and particular bottlenecks that a lack of evaluation principles and guidelines cause – is offered. The contribution of the paper is twofold: a summary of current techniques and approaches for automatic WordNet construction and extension, and a first step in encouraging the further discussion and development of common guidelines for improving the area – particularly in terms of evaluation – moving forward.

2. Background - WordNet Construction

Generally speaking, WordNets are constructed using one of two approaches (Vossen, 1998):

- The *merge* approach – whereby an exhaustive repository of senses (meanings) of each word is compiled, with synsets then created that contain all of the applicable words for a given sense.
- The *expansion* approach – whereby existing synsets from a reference WordNet are used as a guide to create corresponding synsets in a new WordNet, by gathering applicable words that represent the meaning of the synset and ordering them by frequency.

Since the introduction of the PWN and the success of early projects such as *EuroWordNet* that were built around its principles, many projects have focused on building new WordNets in diverse languages using these methods. These endeavours have highlighted various advantages and disadvantages of both the merge and the expansion approaches. Bhattacharyya (2010) describes how the merge approach results in WordNets of high quality, on account of expert lexicographers working in detail on only one language; however, the process is typically very slow. Conversely, the expansion approach can allow the construction of the WordNet to take place much more quickly, with construction

[1] `http://wordnet.princeton.edu/wordnet/`
[2] `http://web.eecs.umich.edu/~mihalcea/downloads.html#semcor`
[3] `http://globalwordnet.org/`

guided by synsets and semantic relationships in the source (or reference) WordNet; however, lexicographers still need to dedicate time to constructing language-specific synsets (meanings or concepts which may not be represented or have a place in the source WordNet), and there is a danger of specific concepts only applicable to the target language being overlooked altogether (Bhattacharyya, 2010).

Years of work on constructing new WordNets have also contributed to the development of guidelines and principles for creating them, largely based on leveraging existing knowledge using the expansion approach. The GWA outlines the importance of 'base concepts'[4] – those concepts that occupy a high position in a semantic hierarchy and have many relations to other concepts – as playing a vital role in constructing WordNets. Base concepts are defined by their *universality* – common (to at least two languages), local (to only one language), or global (across all languages) – with an initial set of 1024 common base concepts being released as part of the *EuroWordNet* project[5]. As a starting point, the GWA proposes that WordNets be constructed in two steps:

- A *core* WordNet of between 5,000 and 10,000 synsets is constructed around the common base concepts,
- An *extended* WordNet is built (semi-automatically, given the semantic basis of the core WordNet) to increase the total number of synsets to 20,000 and beyond.

Given that for many languages these core synsets are readily-constructed, it makes sense to leverage them when constructing new WordNets, and to 'borrow' the semantic relationships that have already been created (Bhattacharyya, 2010). Given the amount of time that can be saved by re-using existing work, there is a tendency to see the expansion approach favoured over the merge approach – it also lends itself extremely well to the automatic construction of synsets, where input from lexicographers is minimal to zero. Thus, the research described in this paper (and particularly in section 3.) largely follows the expansion approach, with synsets being constructed by automatically extracting lexical data from a range of resources in order to build a skeleton framework of meanings based on a reference WordNet.

3. Automatic WordNet Construction from Lexical Resources

Automatic construction and extension of WordNets has traditionally relied on existing lexical resources. Most of the existing research on the subject describes that to create the 'target' WordNet, at least two resources are needed:

- A 'source' WordNet (usually the PWN),
- Lexical resources such as on-line encyclopedias, bilingual dictionaries, or parallel corpora – possibly leveraged in conjunction with additional processing techniques such as machine translation.

[4] http://globalwordnet.org/gwa-base-concepts/

[5] http://www.globalwordnet.org/gwa/ewn_to_bc/topont.htm

3.1. On-line Encyclopedias

An early approach to constructing WordNet synsets automatically explored the idea of extracting semantic relationships from an on-line encyclopedia – namely Wikipedia[6] – in order to extend the PWN (Ruiz-Casado et al., 2005). The approach identifies lexical patterns representing the semantic relationships between entities (links) in Wikipedia, and then uses these patterns to extend existing or create new WordNet synsets. This process takes place over four steps:

- 'Disambiguating' Wikipedia entities to associate them with their corresponding WordNet synsets,
- Extracting patterns of context between Wikipedia entities and other terms they are associated with via a hyperlink,
- Generalising the extracted patterns by comparing them with each other and finding matching ones,
- Applying the patterns to find new semantic relationships not already present in the WordNet.

A manual evaluation of 360 'disambiguated' Wikipedia entities showed that 92% were accurately associated with their corresponding synsets in PWN. As for the patterns extracted from them and used to extend PWN, a total of 1204 hyponomy relations, 418 holonymy relations, and 184 meronymy relations were added to the existing WordNet, with precisions of 0.69, 0.61, and 0.61 respectively. Only 4 hypernymy patterns were extracted by the process (all of which were already present in PWN), perhaps showing that most broader concepts (hypernyms being 'types') are a) likely to be already present in WordNets, and b) not likely to be found in Wikipedia definitions, which by their very nature tend to describe known concepts in more detail.

3.2. Bilingual Dictionaries

The most common technique for populating new WordNets automatically has been to leverage the information in bilingual dictionaries in the source and target languages. Use of bilingual dictionaries for this purpose goes back to very early work on building Catalan and Spanish WordNets as part of the *EuroWordNet* project. In this work, translations of English words in the source WordNet were found, and these translations classified by features such as polysemy (number of translations for each word), structure (the semantic relationships between translations in the source WordNet), and 'conceptual distance' (length of the path between two words in a graph-based representation of the source WordNet) to create a skeleton WordNet in the target language, which could be extended later using bilingual taxonomies (Farreres et al., 1998).

Since then, bilingual dictionaries have continued to be a popular resource for the automatic construction of WordNets. A Romanian WordNet was built by using a range of heuristics to:

- Analyse the relationships between synsets in the source (English) WordNet,

[6] https://en.wikipedia.org

- Identify semantic relationships in various target language resources,
- Map these relationships to each other in the target (Romanian) WordNet using a bilingual dictionary (Barbu and Barbu Mititelu, 2005).

The method was evaluated using 9716 synsets from a pre-existing Romanian WordNet that also had entries in PWN, from which these 9716 synsets were extracted and used as the source (English) WordNet – the synsets used were limited to hypernymy and meronymy relations, and all 19,624 literal words within the synsets had an entry in the bilingual dictionary. The resulting automatically-constructed Romanian WordNet contained 9610 synsets connected by approximately 11,969 semantic relationships, which were reported to be 91% accurate when compared to the 9716 synsets from the pre-existing Romanian WordNet (Barbu and Barbu Mititelu, 2005).

In more recent work on building a Persian WordNet, a bilingual dictionary was used to extract a group of 'candidate' synsets containing English translations of a given Persian word from a source WordNet (PWN). These candidate synsets were then ranked by calculating the Mutual Information of the given Persian word and its English translations in both source and target language corpora, and based on this ranking the most appropriate candidate synset to use for the target (Persian) WordNet was selected (Montazery and Faili, 2010). An extension of this work specifically aimed at lesser-resourced languages was also described, in which a Persian WordNet is constructed by finding the English translations of Persian words in small corpora using a bilingual dictionary. These translations are then used to perform word sense disambiguation (WSD) on a Persian sentence using a source (English) WordNet, and the English synsets returned by the WSD algorithm are mapped to the target (Persian) WordNet (Taghizadeh and Faili, 2016).

Again, these techniques have been shown to be able to automatically construct WordNets with a good degree of accuracy. Montazery and Faili (2010) report that a manual evaluation of 500 synsets from their automatically-constructed target WordNet (which in total covered 29,716 synsets from PWN) resulted in an accuracy of 82.6% (95.8% for synsets whose mapping from source to target WordNet was unambiguous and 76.4% for synsets whose correct mapping had to be decided by ranking multiple candidates). Taghizadeh and Faili (2016) manually evaluated 1,750 word/synset pairs from their target WordNet, and describe how a threshold value (between 0 and 1) used by their WSD algorithm to remove low-scoring candidate synsets had a significant impact on their results. Higher threshold values resulted in the WordNet being more precise (90% with a threshold value of 0.1) but with low recall (fewer synsets in the target WordNet), while lower threshold values resulted in a WordNet with higher recall (more synsets) but with low precision (74% with the threshold value set to 0).

3.3. Machine Translation

While bilingual dictionaries have been the most commonly-used resources in the automatic construction of WordNets, they have also been leveraged in conjunction with additional processing techniques. Recent work by Lam et al. (2014) focused on the construction of WordNets in a variety of languages – Arabic, Assamese, Dimasa, Karbi, and Vietnamese – using machine translation as well as (or in some cases instead of) a bilingual dictionary. For example, they describe:

- An *intermediate WordNets (IW)* approach whereby the same synset from WordNets in a number of different languages is translated into the target language using machine translation,
- An *intermediate WordNets and dictionary (IWND)* approach whereby the same synset from WordNets in a number of different languages is translated into English using a bilingual dictionary, and then from English to the target language using machine translation.

For both techniques, after a set of translated candidate synsets in the target language has been produced, the candidates are ranked based on various heuristics to decide on the most appropriate target language translation for the original synset in question.

Matching subsets of 500 automatically-constructed synsets in Arabic, Assamesse, and Vietnamese were evaluated using a 5-point Likert scale, with an average score of 3.82, 3.78, and 3.75 respectively (3-4, average to good on the Likert scale). It was also shown that the coverage (number of synsets in the automatically-constructed WordNet compared to PWN) of the bilingual dictionary-based IWND technique was generally higher (70,536 synsets in Arabic, for example) than with the IW technique (48,245 synsets with 2 intermediate WordNets, or 61,354 with 3) unless the IW technique was used with 4 intermediate WordNets (75,234 synsets) (Lam et al., 2014).

3.4. Parallel Corpora

A recently-explored alternative to leveraging bilingual dictionaries to automatically construct WordNets has been to exploit parallel corpora in order to map synsets between source and target languages. Oliver and Climent (2014) experimented with extracting synsets for target WordNets by aligning English terms tagged with PWN synset identifiers to their corresponding lemmas in parallel corpora in Spanish, French, German, Italian and Portuguese. This is achieved by assigning synset identifiers from PWN to the English side of the parallel corpora using the *Freeling*[7] and *UKB*[8] word sense disambiguation toolkits, tagging the target language sides with simple POS tags (nouns, verbs, adjectives and adverbs), and then using a simple word alignment algorithm based on most frequent translations to map the subsets of the English synsets to their corresponding source language words.

The resulting automatically-constructed WordNets were evaluated using reference WordNets for each language. For each source language synset for which target language variants were extracted using the method, those target language

[7] `http://nlp.lsi.upc.edu/freeling/node/1`
[8] `http://ixa2.si.ehu.es/ukb/`

variants were compared with the corresponding synset in the target language reference WordNet, with a mapping being considered correct if the target language variant proposed by the method was present in the reference WordNet. Experimenting with three different parallel corpora, the automatically-constructed WordNets were considered quite precise, with results ranging between 75.73% and 85.03% for Spanish, French, Italian and Portuguese (although the extracted German WordNet was noticeably less precise at 45.64-53.15% across the three parallel corpora). However, as with the bilingual dictionary-based approaches to automatic WordNet construction, the reported recall from the parallel corpora-based approach was low (ranging from 10.96 to as low as 2.93 for French synset variants extracted from one of the three corpora). The number of extracted variants in each target language was considered to be few, given the size and number of English PWN synsets present in the English sides of the corpora.

4. Word Embeddings and WordNet Synsets

Given the increasing popularity of word embeddings (vector space representations of word meanings based on their distribution within large datasets), it should come as no surprise that the links between embeddings and traditional, WordNet-style representations of word senses have recently been explored. Nayak (2015) demonstrated that word embeddings can be classified according to words and can also be used to predict hypernymy relations between them, while Rothe and Schütze (2015) report that sets of embeddings trained not just on words but also on synsets (groups of synonyms) and lexemes (word-synset pairs) achieve state-of-the-art performance on WSD and semantic similarity tasks. This kind of research shows the potential of word embeddings for capturing the kinds of relationships (and being useful in the types of tasks) commonly associated with WordNet-style word senses – potential which is further compounded by the reported high precision with which multi-sense word embeddings can be mapped to WordNet-style synset entries in Babelnet[9] (Panchenko, 2016).

4.1. Extending and Constructing Synsets from Word Embeddings

Naturally, recent research has explored leveraging the links between word embeddings and synsets in order to automatically construct new synsets from the embeddings themselves. Sand et al. (2017) describe using word embeddings to extend an existing Norwegian WordNet by finding candidate hypernyms for a given word based on its nearest neighbours in the WordNet, and then scoring these candidate hypernyms by distributional similarity (using the vector space of the embeddings) and distance in a graph-based representation of the WordNet. Based on an evaluation of 1388 target words occurring 5 times or more in the news corpus on which the embeddings were trained, an *accuracy* (percentage of newly-added target words correctly placed under the appropriate hypernym) of 55.80% and an *attachment* score (percentage of target words actually added to the Norwegian WordNet) of 96.33% were recorded. This accuracy is increased when only evaluating on target words that occurred more than 100, or more than 500 times in the corpus, but at the cost of diminished coverage (fewer target words available with which to extend the Norwegian WordNet).

An alternative approach to using word embeddings to extend an existing WordNet has been described by Al tarouti and Kalita (2016), who in fact use word embeddings to extend an automatically-constructed Arabic WordNet built using the machine translation / bilingual dictionary method described by Lam et al. (2014). They leverage word embeddings to compute the cosine similarity of a) words within candidate synsets, and b) words within pairs of semantically-related synsets, allowing them to discard candidate synsets (and words within them) whose cosine similarity is below a given threshold value. 600 automatically-constructed word pairs (of synonym, hypernym, holonym, and meronym types) were evaluated by Arabic speakers using a 5-point Likert scale, with the average score then converted to a percentage – the resulting precision of the synonyms, hypernyms, holonyms, and meronyms was 78.4%, 84.4%, 90.4%, and 79.6% respectively, slightly higher than the precision (as a percentage) reported by Lam et al. (2014) for Arabic.

A similar method has also been described by Khodak et al. (2017), who report on the automatic construction of whole WordNets in French and Russian from scratch using bilingual dictionaries and word embeddings. After producing candidate synsets by finding the corresponding source language synsets for target language words as given by the bilingual dictionaries, word sense embeddings and word sense induction (WSI) techniques are used to cluster only the most relevant translations of lemmas from the source language synset together, ensuring that the correct target language candidate synset is 'matched' as correct. Evaluating these methods using subsets of 200 nouns, verbs and adjectives from each of the target language WordNets, the resulting F_5 scores – used as a precision-centric alternative to the usual F_1 score – were reported to outperform those yielded using a baseline similarity method by 6% and 10% for French and Russian respectively.

5. Issues for Evaluating Automatically-Constructed WordNets

One of the biggest issues for the automatic construction of WordNets is how to properly and effectively evaluate their accuracy and/or precision. Across all of the different lexical resource and word embedding-based approaches to automatic synset construction described in the previous two sections, evaluation methods can be split across two types:

- Comparison against a reference WordNet,
- Manual evaluations against fixed samples of automatically constructed synsets.

Focusing first on comparisons with reference WordNets, much of the research referenced in the preceding sections reports on problems with this kind of evaluation. Khodak et al. (2017) describe an attempted comparison with their

[9] http://babelnet.org/

automatically-constructed WordNets and reference WordNets being difficult, with the ELRA French WordNet 2 being only around half the size of their new French WordNet and most Russian WordNets being a) even smaller, and b) not easily linked to (or compared with) PWN. Similarly, Taghizadeh and Faili (2016) cite 'the lack of correct links' in the pre-existing *FarsNet* (an ontology of Persian words mapped to PWN synsets) as being troublesome when attempting to compare their automatically-constructed Persian WordNet to it – they reported after comparing their automatically-constructed Persian WordNet to *Farsnet* that that the precision of their new WordNet was just 19% and its recall 49%, too low 'to be considered as a reliable resource'.

The discrepancies between size and coverage of reference WordNets and the original PWN can be viewed as an issue of granularity: PWN is large enough that its senses are fine-grained, and so several PWN synsets can generally mapped onto one synset in a reference WordNet (such as *FarsNet*) while other PWN synsets will not be present at all (Khodak et al., 2017). This makes it difficult, when comparing automatically-constructed WordNets to reference WordNets in a target language, to decide whether newly-created synsets are correct or not. For example, Taghizadeh and Faili (2016) use the following criteria to decide whether word and synset pairs in an automatically-constructed Persian WordNet are correct, using *FarsNet* as their reference WordNet:

- If a Persian word does not exist in *FarsNet*, it *IS NOT* correct,
- If a Persian word exists in *FarsNet* but is not linked to a PWN synset, it *IS NOT* correct,
- If a Persian word exists in *Farsnet* and and at least one PWN synset is linked to it:
  - If the automatically-constructed synset is not one of the linked PWN synsets in *FarsNet*, it *IS NOT* correct,
  - IF the automatically-constructed synset is one of the linked PWN sysnets in *FarsNet*, it *IS* correct.

Out of three options here, two of them lead to the word in the automatically constructed Persian WordNet being classed as incorrect – and even if the word is in both *FarsNet* and PWN, that word still has to be linked between *those* resources to be accepted as correct. This approach is therefore totally dependent on the quality of *FarsNet*, and any words in the automatically constructed WordNet that are in PWN and that *should* be in *FarsNet* will, unfortunately, be classed as incorrect. As Oliver and Climent (2014) – who considered an automatically extracted synset correct only if it was also present in a reference WordNet – highlight, automatic comparisons with reference WordNets inevitably mean that if the reference WordNets are not complete, then correctly extracted synsets in the automatically constructed WordNet can be evaluated as incorrect – and this is a major problem when reporting on their accuracy and legitimacy as a lexical resource.

Sand et al. (2017) also touch on potential discrepancies between automatically extracted synsets and their equivalent synsets in reference WordNets or in PWN, noting that hypernymy relations 'can be right or wrong by varying degrees'. They describe a 'soft accuracy' measure whereby the accuracy of an automatically extracted synset is weighted according to the number of links (or edges) between words in different synsets that separate a given word from what would be its correct position in a graph-based representation of the WordNet. Weighting the accuracy of automatically extracted synsets according to how comparable they are with their fully-formed PWN equivalents is certainly more logical than evaluating strictly on 'correct insertions' – an automatically extracted synset containing 8 of the 10 links to other words present in the same sysnet in PWN, for example, is surely more correct than an automatically extracted synset containing only 2 or 3 of the 10 links.

The alternative to automatic evaluations of synset correctness is of course manual evaluation, which is widely used both in isolation and in conjunction with automated evaluations in the works cited in Section 3. (Ruiz-Casado et al., 2005; Montazery and Faili, 2010; Lam et al., 2014; Taghizadeh and Faili, 2016). However, as Ruiz-Casado et al. (2005) point out, it is 'difficult to know how accurate manually-evaluated synsets are without some common guidelines. Some works simply describe having manual annotators decide if an automatically extracted is or is not semantically similar to a reference synset (Taghizadeh and Faili, 2016), while others – much more in line with the idea of weighting accuracy according to a degree of correctness (Sand et al., 2017) – have used a Likert scale for conducting manual evaluations (Lam et al., 2014).

6. Conclusions

This paper has examined a number of approaches to automatically construct and extend WordNets, taking into account both lexical and word embedding-based approaches to extracting synsets. A common trend across evaluations of WordNets extracted from lexical resources is that while the synsets themselves are reasonably precise, recall is often very low – that is, although extracted synsets are accurate when compared to reference WordNets such as PWN, not enough synsets are actually being extracted using automatic methods. However, these results do not necessarily paint the full picture – there are few agreed principles or common guidelines for evaluating extracted synsets, and it is often difficult to decide what constitutes a correctly or incorrectly extracted synset.

Automatic WordNet construction is a promising research area, particularly in the context of lesser-resourced languages – the fact that the works outlined in Sections 3. and 4. cover typically under-resourced languages such as Arabic, Catalan, Persian (Farsi), Romanian and Russian demonstrates that researchers see the value in exploring how to improve it. Constructing WordNets manually takes many years of linguistic knowledge, making them a costly and labour-intensive endeavour, but the availability of lexical resources in many languages and widespread adoption of approaches such as word embeddings for extracting meaning and relationships from free text have made the concept of automatically constructing accurate WordNets in lesser-resourced languages very feasible.

This paper is intended to provide a summary of commonplace techniques in automatic WordNet construction and extension. In highlighting some of the underlying issues that have been a bottleneck for evaluating them, it is also hoped that it can serve as a first step in encouraging further discussion and the development of common guidelines for improving the area moving forward. A set of agreed principles that help researchers paint a clearer picture of what constitutes a correctly or incorrectly extracted synset will be an important next step in encouraging the automatic construction and extension of WordNets, particularly for those working with lesser-resourced languages.

References

Al tarouti, F. and Kalita, J. (2016). Enhancing Automatic WordNet Construction Using Word Embeddings. In *Proceedings of the Workshop on Multilingual and Cross-lingual Methods in NLP*, pages 30–34, San Diego, California. Association for Computational Linguistics.

Barbu, E. and Barbu Mititelu, V. (2005). Automatic building of Wordnets. In *Proceedings of International Conference on Recent Advances in Natural Language Processing*, RANLP 2005.

Bhattacharyya, P. (2010). IndoWordnet. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, LREC '10. European Language Resources Association.

Farreres, X., Rigau, G., and Rodríguez, H. (1998). Using WordNet for Building WordNets. In *Proceedings of the COLING-ACL '98 Workshop on Usage of WordNet in Natural Language Processing Systems*.

Fellbaum, C. (1998). *WordNet: An Electronic Lexical Database*. MIT Press.

Khodak, M., Risteski, A., Fellbaum, C., and Arora, S. (2017). Automated wordnet construction using word embeddings. In *Proceedings of the 1st Workshop on Sense, Concept and Entity Representations and their Applications*, pages 12–23, Valencia, Spain. Association for Computational Linguistics.

Lam, K. N., Al Tarouti, F., and Kalita, J. (2014). Automatically Constructing Wordnet Synsets. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Short Papers)*, ACL 2014, pages 106–111. Association for Computational Linguistics.

Montazery, M. and Faili, H. (2010). Automatic Persian WordNet Construction. In *Proceedings of the 23rd International Conference on Computational Linguistics*, COLING '10, pages 846–850. Association for Computational Linguistics.

Nayak, N. (2015). Learning hypernymy over word embeddings. Technical report, Stanford University.

Oliver, A. and Climent, S. (2014). Automatic Creation of WordNets from Parallel Corpora. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation*, LREC '14. European Language Resources Association.

Panchenko, A. (2016). Best of Both Worlds: Making Word Sense Embeddings Interpretable. In *Proceedings of the 10th Language Resources and Evaluation Conference*, pages 2649–2655, Portorož, Slovenia. European Language Resources Association.

Pianta, E., Bentivogli, L., and Girardi, C. (2002). MultiWordNet: Developing an Aligned Multilingual Database. In *Proceedings of the First International Global WordNet Conference*.

Rothe, S. and Schütze, H. (2015). Autoextend: Extending word embeddings to embeddings for synsets and lexemes. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, ACL 2015, pages 1793–1803, Beijing, China. Association for Computational Linguistics.

Ruiz-Casado, M., Alfonseca, E., and Castells, P. (2005). Automatic Extraction of Semantic Relationships for WordNet by Means of Pattern Learning from Wikipedia. In *Proceedings of the 10th International Conference on Natural Language Processing and Information Systems*, NLDB '05, pages 67–79. Association for Computational Linguistics.

Sand, H., Velldal, E., and Øvrelid, L. (2017). WordNet Extension via Word Embeddings: Experiments on the Norwegian WordNet. In *Proceedings of the 21st Nordic Conference of Computational Linguistics*, pages 298–302, Gothenburg, Sweden. Linköping University Electronic Press.

Taghizadeh, N. and Faili, H. (2016). Automatic Wordnet Development for Low-resource Languages Using Cross-lingual WSD. *Journal of Artificial Intelligence Research*, 56(1):61–87.

Vossen, P. (1998). *EuroWordNet: A Multilingual Database with Lexical Semantic Networks*. Springer Netherlands.

Vossen, P. (2004). EuroWordNet: A Multilingual Database of Autonomous and Language-Specific WordNets Connected via an Inter-Lingual Index. *International Journal of Lexicography*, 17(2):161–173.