

Multilingual Wordnet sense Ranking using nearest context

Abstract

In this paper, we combine methods to estimate sense rankings from raw text with recent work on word embeddings to provide sense ranking estimates for the entries in the Open Multilingual Wordnet (OMW). The existing Word2Vec Polyglot2 pre-trained models are only built for single word entries, we, therefore, re-train them with multiword expressions from the wordnets, so that multiword expressions can also be ranked. Thus this trained model gives embeddings for both single words and multiwords. The resulting lexicon gives a strong WSD in five languages. The results are evaluated for Semcor sense corpora for 5 languages using Word2Vec and Glove model. The Glove model achieves the average accuracy of 0.47 and Word2Vec achieves 0.31 for languages such as English, Italian, Indonesian, Chinese and Japanese. The experimentation on OMW sense ranking proves that the rank correlation mostly similar to the human ranking. Hence distributional semantics certainly aid in Wordnet Sense Ranking.

1 Introduction

Most of the existing Wordnet sense rankings (Navigli, 2009) use document level statistics to find the prominent sense of the given word. (McCarthy and Carroll, 2003) showed that predominate senses could be learned from a sufficiently large corpus, and this work has since been extended by various researchers. (Lim, 2014; Liu et al., 2015; Pocostales, 2016; Rong, 2014; Long et al.,

2016) Words that appear nearest to the given word conveys the context/meaning of a word, gives the most frequently used senses of a given word. This proposed work uses nearest context words to predict the senses and computes the frequency of occurrence of these senses within the corpus. Since most of the existing WSD systems utilize the Most Frequent Sense (MFS) as a baseline, it is important to rank the Wordnet senses in a meaningful way. Many Word embedding techniques are based on n-gram models and unsupervised learning (Bhingardive et al., 2015a). However recent word embeddings are based on the neural network architecture that can train the models faster and more accurately than traditional approaches.

Two well-known software packages used to train word embeddings, are Word2Vec (Rong, 2014) and Glove model (Pennington et al., 2014). Deep learning approaches for word embedding (Tang et al., 2014) are recent research area that learns multiple levels of features to handle complexity. Naturally, every deep learning neural network takes words from a vocabulary and embeds them as a low dimensionality vectors and fine-tunes through back-propagation in the subsequent layers. The main difference between such a network that produces word embeddings with the word2vec is their computational complexity. Generating word embeddings with a very deep architecture is highly expensive for a large corpus. Another difference is that semantically coherent embedding can be obtained with the help of deep learning based approaches. Some optimization techniques (Pennington et al., 2014) are required to reduce the time and computational complexities of these models. Polyglot (Al-Rfou et al., 2013) is a natural language pipeline that supports many

NLP based tasks such as tokenization, Language detection, Named Entity Recognition, Part of Speech Tagging, Sentiment Analysis, Word Embeddings, Morphological analysis and Transliteration for many languages. This work utilizes their Word embeddings. Existing polyglot word embeddings (Al-Rfou et al., 2013) support 137 languages. We have planned to use the word embeddings for the 35 hand-built wordnets currently in OMW (Ruci, 2008; Elkateb et al., 2006; Borin et al., 2013; Pedersen et al., 2009; Simov and Osenova, 2010; Gonzalez-Agirre et al., 2012; Pociello et al., 2011; Wang and Bond, 2013; Huang et al., 2010; Pedersen et al., 2009; Fellbaum, 1998; Stamou et al., 2004; Sagot and Fišer, 2008; Ordan and Wintner, 2007; Mohamed Noor et al., 2011; Isahara et al., 2008; Montazery and Faili, 2010; Lindén and Carlson., 2010; Garabík and Pileckytė, 2013; Vossen and Postma, 2014; Piasecki et al., 2009; de Paiva et al., 2012; Tufiş et al., 2008; Darja et al., 2012; Borin et al., 2013; Thoongsup et al., 2009; Pianta et al., 2002; Oliver et al., 2015; Raffaelli et al., 2008; Toral et al., 2010).

We use corpus based frequencies for five of these languages (English, Chinese, Japanese, Italian and Indonesian) from the NTU Multilingual Corpus (Tan and Bond, 2013) and use them to evaluate the learned sense rankings. Our major contribution is training and testing on large numbers of multiword expressions, which are often neglected in the word embedding literature. We identify the multiword expressions identified in the hand-built lexicons and train our own model for them using Word2Vec (Rong, 2014) and Glove (Pennington et al., 2014).

This paper is structured as follows. Section 2 discusses the related work in Word embedding and its application in Wordnet Synset Ranking. Section 3 describes the data, methods, and Section 4 discusses the evaluation of results obtained from word embedding and its effect in Wordnet Sense Ranking. Finally, Section 5 concludes with the findings and future plans to improve the results.

2 Related Work

Word embedding techniques have been popular in recent years in Word Sense Disambigua-

tion (*WSD*) research. similar to this proposed work, (Bhingardive et al., 2015b) computes word embeddings with the help of pretrained Word2Vec and matches with the sense embeddings obtained from the Wordnet features. They have attempted Wordnet sense ranking for Hindi and English. Since the Word2Vec model is based on the frequency of occurrence of words in the corpus, finding the nearest context words that occur infrequently in the corpus is difficult.

(Panchenko, 2016) compares sense embeddings of AdaGram (Bartunov et al., 2015) with BabelNet (Navigli and Ponzetto, 2010) synsets and proved that sense embeddings can be retrieved by automatically learned sense vectors. Sense embeddings for a given target word are identified by finding the similarity between the AdaGram Word embeddings list with the BabelNet Synsets words list. (Rothe and Schütze, 2015) proposed an approach that takes word embeddings as input and produces synset, lexeme embeddings without retraining them. They used Wordnet lexical resource to improve word embeddings.

(Arora et al., 2016) showed that word vectors can capture polysemy and word vectors can be thought of as linear superpositions of each sense vector. They have attempted discourse analysis to find the cluster of sense vectors.

Although the basic idea of word embeddings is not tied to any one languages, the preprocessing steps are language specific. (Kang et al., 2016) presented a cross-lingual word embedding for English and Chinese Word Sense Disambiguation (*WSD*). They have experimented with the performance of *WSD* using different word embeddings such as Word2Vec and Glove model. (Bhingardive et al., 2015a) compared word embeddings obtained from the Word2Vec model and the sense embedding obtained from the Wordnet for English and Hindi languages and restricted to Nouns. They used various Wordnet features similar to this proposed work to find the predominant sense. Their approach outperforms SemCor baseline for words with the frequency below five.

Sentence embedding is an another interesting research area mostly applied in machine translation systems. (Palangi et al., 2016) have

proposed sentence embedding using recurrent neural networks (RNN) with Long Short-Term Memory (LSTM) cells. The main objective of the learning model is to compute a weight factor that makes semantically similar sentences as close as possible and semantically dissimilar sentences as far as possible. This helps us to learn the sentence context not only based on the nearest context but also their semantic similarity. In sentence embedding, semantically similar sentences that occur close to the given sentence are obtained. Most of the sentence embedding approaches use deep learning architectures to reduce the time required to train the model. Recurrent neural networks (RNN) with Long Short-Term Memory (LSTM) cells have been used due to its ability to capture and accumulate long-term information about the previous sentences. Though the sentence embedding is computationally costly, we can apply this approach to further improve the wordnet sense embedding.

Another approach to finding predominant senses for multiple languages used the k-Nearest-Neighbors (k-NN) algorithm to find Distributional Semantic Similarity (DSS) for the word-pair (Lim, 2014). They used various similarity measures such as Wu-Palmer, Least Common Subsumer (LCS) and Leacock-Chodorow (Pedersen et al., 2004) to measure the distance between two senses. In this research context words are identified with the help of Polyglot (Al-Rfou et al., 2013) word embeddings.

3 Methodology

In this work, we use word embeddings to find the nearest context of a given word and compare it with the senses obtained from the OMW to find the most frequently used senses. Our aim is to rank the senses obtained from the OMW with the help of the context words and their frequency of occurrence. Initially, we use the pretrained polyglot word embedding model (Al-Rfou et al., 2013) to retrieve the nearest context words and found multiwords are unidentified. Hence in this work, we have trained our own model similar to polyglot for both single and multi-words. Our aim is to train this model for all 35 languages supported

by OMW, for this paper we present only the results for the five languages for which we have evaluation data: English, Chinese, Japanese, Italian and Indonesian.

3.1 Corpus Cleaning and preprocessing

We exploit the openly available Polyglot wiki dump corpus (Al-Rfou et al., 2013) for English, Chinese, Japanese, Italian and Indonesian. Since polyglot wiki dump corpus supports various domains and languages, we have used this corpus to train our own model. Before training our own model, the corpus texts are preprocessed by removing symbols, numbers and shortest text. Stop words have been removed with the help of the NLTK toolkit (Bird et al., 2009). However, NLTK does not support stopwords for all languages. Hence we have included stop words of Chinese, Japanese, Indonesian, Italian from publicly available online utilities to NLTK toolkit. For English, Indonesian and Italian we have lemmatized each word of the cleaned text to find their base form. Chinese does not inflect, and Japanese inflections are normally split off by the tokenizer. Hence we have used Mecab to tokenize/lemmatize Japanese texts. After preprocessing the text, each sentence of the corpus is tokenized into single and multiple terms. In order to identify the multiwords from the corpus, we have used the existing Wordnet MWE lexicon (*MWEs*). The terms of sentences are matched with the existing wordnet *MWEs* lexicon and merged the Multiwords tokens with ”_” symbol if the terms are available in the *MWE* lexicon. This preprocessed MWE tagged texts are given as input to train our own model. So, for example, a sentence like *I looked five words up* will be preprocessed to `I look_up word`.

3.2 Training Model

Word embeddings for the above five languages have been trained using the Polyglot2 (Al-Rfou et al., 2013) package and *Global Vectors for Word Representation Glove Model* (Pennington et al., 2014). Polyglot2 is a software package that enables building your own language models. It learns the distributed representations of words/word embeddings for the given corpus. The glove is another unsu-

pervised learning algorithm used for obtaining vector representations of words. Training is performed by considering global word-word co-occurrence statistics from a corpus and results with the linear substructures of the word vector space. We can build our own word embeddings with the help of Polyglot2 and Glove models.

3.3 Predominant Sense Scoring

To find the predominant senses for the given word w , the senses obtained from the OMW are represented as $S_w = S_1, S_2, \dots, S_n$. The neighbouring context obtained from Polyglot2 or Glove is represented as $S_w^N(w, d)$ where N represents the number of neighbouring contexts from word embedding obtained for the senses S_w that can vary from 1 to N , and d represent the distance score between the S_w and S_w^N . $P_s(S_w)$ represents the predominant score of S_w based on the wordnet synset similarity.

$$P_s(S_w) = \log(\text{sum}(S_w^N(w, d)) * M_W^T / T^N W_e) \quad (1)$$

M_W^T - represents the number of matching terms between the *OMW* synset definitions and example sentences with respect to polyglot word embeddings

$T^N W_e$ - represents the number of word embeddings obtained from Polyglot2.

After computing the predominant score $P_s(S_w)$ for each wordnet entries the semantic similarity between the word embedding with the OMW ontology hierarchy (S_H) is measured.

$$S_H = P_s(S_w) * H_s(M) / T^N W_e \quad (2)$$

$H_s(M)$ represents the number of concepts such as Hypernyms and Hyponyms of WordNet Ontology that match with the number of terms obtained in the polyglot word embeddings. The intuition behind is that the words in wordembedding will have similar words that can appear in WordNet hierarchy. For example, the word *party* may refer to a person, organization or an occasion. If it refers to a *person*, the hypernyms are *person* and the hyponyms are *assignee*, *assignor*, *contractor*, *intervenor*. Similarly for organization the hypernyms is *set* and hyponyms are *fatigue_party*,

landing_party, *party_to_the_action*, *rescue_party*, *search_party*, *stretcher_party*, *war_party* and for considering *occasion* as sense the hypernyms are *affair* and hyponyms are *bash*, *birthday party*, *bunfight*, *ceilidh*, *cocktail_party*, *dance*, *fete*, *house_party*, *jolly*, *tea_party*, *whist_drive*.

When we give *Person* as Input to Polyglot2(Al-Rfou et al., 2013), we will get the following word embeddings. *person-0.575121*, *contractor-0.628679*, *team-0.619203*, *division-0.682174*, *unit-0.700489*, *government-0.62491*, *strategy-0.725378*, *event-0.692839*, *camp-0.689145* *program-0.688767*. The terms such as *person* and *contractor* matched with the Wordnet hypernyms and hyponyms. Thus *person* sense is the most predominantly used when compared to *organization* and *event* senses since it shares the semantics with WordNet hierarchy. Similarly, we can match with other features of WordNet senses to infer which sense is important.

4 Results and Evaluation

In this section, the word embedding models such as Glove(Al-Rfou et al., 2013) and Polyglot2(Word2Vec)(Pennington et al., 2014) have been evaluated on two different tasks such as word-sense ranking of Wordnet and query expansion for clinical texts, then we present some examples of word embeddings for intuitive comprehension. The word sense ranking and trained word embeddings have been tested for 5 languages English, Chinese, Japanese, Indonesian and Italian languages of Semcor dataset for the words with more than one sense. In order to train the multi-words in Polyglot embedding model(Al-Rfou et al., 2013), we have assigned the Context Window Size as 14, Initial learning rate as 0.025, Hidden Layer size as 32 and minimum word count as 2. In Glove word embedding, we have set the minimum word count as 2, Vector size as 100, Maximum Iteration as 100 and Context Window size as 14.

We have measured the accuracy of word embeddings based on the human relevant judgment. Among the top 10 word embeddings, the fraction of relevant embeddings is measured. The results are shown in Table 1 and Table 2. We find that the trained Glove

Languages	Accuracy
English	0.65
Chinese	0.42
Japanese	0.34
Italian	0.41
Indonesian	0.54

Table 1: Accuracy of Word embedding score for Glove (Trained) Model

Languages	Accuracy
English	0.32
Chinese	0.22
Japanese	0.14
Italian	0.32
Indonesian	0.24

Table 2: Accuracy of Word embedding score for Polyglot2 (Trained) Model

model gave a better result when compared to the *Polyglot2(Word2Vec)*(Al-Rfou et al., 2013) model. Even though Word2Vec trained the model using Polyglot2 ¹ gave a less relevant result, we found that the existing pretrained model(Al-Rfou et al., 2013) gave better results than our trained model. In order to test this, We have taken 5611 unique terms and found that existing pre-trained model handles 1500 terms semantically correct and the remaining 4111 terms are not handled. The reason is pre-trained polyglot2 Word2Vec model is trained on wiki corpus and unable to scale up to the medical domain. Moreover, it is not trained for Multiwords.

Few sample list of analyzed words are given below.

List of semantic-based word embedding obtained in each language(Polyglot2) are listed below: Indonesian

- *lokasi(location)*
:wilayah,peta,batas,daerah,stasiun,jalur ,pelabuhan,kawasan,tujuan,arah
(region, map, boundary, area, station, line
ports, regions, destinations, directions)

Italian:

- *paura(fear)* - *frustrazione,sensazione,sofferenza*
(fear - frustration, feeling, suffering)

English:

- *Patients- drugs, medications, sufferers, radicals, physicians, agents, genes, individuals, clients, viruses*

Japanese:

- *条約 (Treaty)*
: 総督, 協定, 紛争, 裁判所, 議会, 法典, 憲章

- *Governor, agreement, conflict, court, parliament, law, charter*

Chinese:

- *网络*
: 电脑, 软件, 计算机, 技术, 电子

(*Network*

: *computer, software, computer, technology, electronics*

)

Since we have trained the model for multi-word expressions, we have specifically analyzed the embeddings for multi-words and the resultant samples are shown below.

Sample List of Multi-words and Nearest Context Word:

- *Query—English:*
deficit_hyperactivity_disorder:
– *attention, memory,deficit_hyperactivity_disorder, adhd,rigidly, proliferative, splinted, treat_attention, allergic_rhinitis, special*
- *Query—Japanese:*
プリンス _ オヴ _ ウェールズ (Prince of Wales):
– *トレハラーゼ, ろかく, レゼルヴ, フリーア, グローヴス, レインボーカップファイナル,mishnaic,traininfomation, カタリココ*
– *(Trehalase, fighting, reserve, free, Groves, Rainbow Cup Final, mishnaic,traininfomation, Catalina Coco)*

¹<https://sites.google.com/site/rmyeid/projects/polyglot>

- *Query– Chinese:*
足球 _ 运动员 (soccer player):
 - 大 _ 祭台, 阅览, 鑑, 諫, 分内事, 大捷, 新交, 續, 井底
 - (*Large altar, learning clang, remonstrance, sub-ministry, victory, new cross, play, bottom*)
- *Query– Indonesian:*
erosi _ pantai (erosion):
 - : *Mikrokimerisme, gerunggang, membuat_bangkrut, mikkeli, lille, superintendent, thur, cibinuang, operasi_boolean*
 - (*Mikrokimerisme, gerunggang, membuat_bangkrut, mikkeli lille, superintendent, thur, ibinuang, operasi_boolean*)
- *Query– Italian:*
seconda_guerra_mondiale (Second World War):
 - *tisiddu, smetliviyi, pelligra, mortificava, skavronskij, tureaud, preprocessing, telemolise, quetzalctl*
 - (*Mixed with other language text*)

Results of semantic based word embedding obtained for each language of Glove are listed below:

- *Seconda_guerra_mondiale (Second World War) (Italian):*
 - *prima_guerra_mondiale, scoppio, guerra, conflitto, dopoguerra, militare, bellico, militari, guerra_mondiale, sovietica* (WWI, outbreak, war, conflict, war, military, war, Word war, military)
- *jus_lemon (lemon juice) (Indonesian):*
 - *Memberikan_tenaga, Mengasamkan, operated, menguapkan, boya, memfermentasi, efektif, recoil, mwh, meluapkan. (provide power, acidity, operated, Evaporate, boya, ferment, effective, recoil, mwh, vent)*

Accuracy(Word2Vec)	Accuracy(Glove)
0.35	0.67

Table 3: Accuracy of Word embedding score for medical text(English)

- *Chinese: 参考 _ 资料 (Reference_Information):*
 - 注释, 脚注, 参考, 迈, 资料 _ 来源, 内部 _ 网络, 注解, 服务 _ 设施, 参见, 出处 (*Reference _ information, Annotations, Footnote, reference, Side, Information source, Internal _ network, annotation, Service Facilities, See also, Source*)
- *English: Treadmill_test:*
 - *Stress_test, exercise, physiology, suggestion, participate, vigorous, walking, prescription, intensity*

We are able to get semantically relevant results even for multi-word expressions with the help of Glove model. Sample results obtained from each language are shown below. We found relevant results in English, Chinese, Italian; However for Indonesian documents, the results are often mixed with other language texts, even though we are able to get meaningful word embeddings. We also found that the Japanese text corpus is tagged with minimal multi-word expressions and noisy. The reason is Japanese text has different writing styles that degrade the accuracy of MWE tagging because the MWE lexicon basically includes the standard scripts. Hence we need to fine tune the MWE tagging by properly filtering the character-level, word level non-standard noisy text.

In order to check, the scalability of these models in different domains, We have tested with Singapore Clinical Practical Guidelines documents of Dental, Medical, Nursing, and Pharmacy of 72 documents, available from *Ministry of Health*, Singapore (2016).² There are 124.2 MB in all. The results are shown in Table 3.

Again the accuracy of Glove model³ is better when compared to the Word2Vec Polyglot

²They are online at https://www.moh.gov.sg/content/moh_web/healthprofessionalsportal/doctors/guidelines/cpg_medical.html.

³<https://nlp.stanford.edu/projects/glove/>

learned model because Glove model computes co-occurrence statistics across the corpus whereas Word2Vec computes co-occurrence statistics within the context window size. The word embedding results also depend on the context window size and minimum frequency count. If we increase both the context window size and minimum frequency count to a certain extent, we can achieve semantically relevant word embeddings. However, the recall will be low.

In order to find the optimum value to maintain precision and recall, we need to run the test with different values for few test samples. The quality and size of the corpus may also impact the results. Since clinical text contains only domain-specific terms which are unambiguous, we are able to achieve meaningful results. Whereas We found difficulty in Wikipedia dump corpus(5 languages) because it contains a lot of noisy data. Our purpose of this work is to check, how far this distributional semantics can help in Word Sense Ranking and Clinical Information Retrieval.

Since existing polyglot model (Al-Rfou et al., 2013) handles single terms well and the trained glove model (Pennington et al., 2014) handle most of the terms meaningfully we have planned to merge both the models to handle single and multi-terms word embeddings.

To evaluate the quality of rankings produced by this method, we have compared the OMW search interface random ranking (*Approach 1*) *A1*, Word embedding (*Approach 2*) *A2* and human judgment rank (*Approach 3*) *A3* list. There are basically two well-defined algorithms such as *Spearman's* and *Kendall's tau* (Kumar and Vassilvitskii, 2010) rank correlation have been used to find the statistical difference in ranking. DCG (*Discounted Cumulative Gain*)(Harman, 2011) measures both relevance and ranking, whereas rank correlation helps to find statistically significant difference in order. *Webber et al (2010)* (Webber et al., 2010) proposed a method to compare ranking quality of two methods and addressed the top-relatedness issue. Since this proposed work needs to consider the concordance and discordance of ranked results based on position, We have used this measure to find the correlation between the two ranked lists. The

Synsets (gleam)

be shiny, as if wet
 shine brightly, like a star or a light
 appear briefly
 an appearance of reflected light
 a flash of light (especially reflected light)

Table 4: Ranking result of Approach 2

Synsets (gleam)

be shiny, as if wet
a flash of light (especially reflected light)
shine brightly, like a star or a light
appear briefly
an appearance of reflected light

Table 5: Ranking result of Approach 1

correlation score is measured with *Approach 2* to *Approach 1* and *Approach 3* for the terms in Semcor dataset. The results are shown in table 7.

For example, when we give "gleam" as query, the resulted ranking of *A1*, *A2*, *A3* are shown in Table 5, Table 4, Table 6 respectively. The rank overlapping between Approach 2 to Approach 1 and Approach 2 to Approach 3 are calculated. Here in this example, the baseline random ranking (Approach 1) is dissimilar from the second position onwards, whereas with human judgment (Approach 3) only the 3rd synset is moved to the last position and the remaining ranking is similar to the proposed approach. Hence the Rank correlation for Approach 2 to Approach 1 is 0.78 and Approach 2 to Approach 3 is 0.88. Thus the rank quality depends on how much it is similar to the human judgment. There are other set-based approaches are also used to compare two ranked lists. However, these approaches require additional computation in handling disjointness and top-relatedness issue when comparing the ranked list with dissimilar content where the positions are important. In this work, the ranked contents are similar and relevant in all the three approaches, but the position

Synsets (gleam)

be shiny, as if wet
 shine brightly, like a star or a light
 an appearance of reflected light
 a flash of light (especially reflected light)
appear briefly

Table 6: Ranking result of Approach 3

Languages	A1 to A2	A2 to A3
English	0.55	0.75
Chinese	0.62	0.68
Japanese	0.64	0.69
Italian	0.61	0.67
Indonesian	0.44	0.56

Table 7: Average Rank correlation analysis between A1 to A2 and A2 to A3

of concordance and discordance are important.

We found that the average correlation between A2 to A3 is greater than A1 to A2. This result provides an additional validation of our model as it demonstrates that the sense ranking can capture the sense preferred by a human. Hence the word embedding score definitely aid in wordnet sense ranking.

5 Conclusion

OMW has over 150 languages with word-nets built automatically, ranging from major languages like German or Korean for which there are no free word-nets, to smaller languages such as Volapuk. For all languages for which Polyglot has data (which is most of them) we will learn rankings and incorporate them into OMW, so that the lexicon is maximally useful for speakers of as many languages as possible.

In future, we planned to extend this work to identifying missing senses by comparing the trained model over the sense-annotated corpus with the existing pre-trained models like polyglot. Since the Glove model is based on co-occurrence context, it gave better results even for a tiny corpus, hence we have planned to extend our model to sentence embedding using Glove model for finding nearest context sentences for a given synset example sentence to further improve our wordnet ranking.

References

Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2013. Polyglot: Distributed word representations for multilingual nlp. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 183–192, Sofia, Bulgaria, August. Association for Computational Linguistics.

Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. 2016. Linear algebraic

structure of word senses, with applications to polysemy. *arXiv preprint arXiv:1601.03764*.

Sergey Bartunov, Dmitry Kondrashkin, Anton Osokin, and Dmitry Vetrov. 2015. Breaking sticks and ambiguities with adaptive skip-gram. *arXiv preprint arXiv:1502.07257*, pages 47–54.

Sudha Bhingardive, Dharendra Singh, Rudra Murthy, Hanumant Redkar, and Pushpak Bhattacharyya. 2015a. Unsupervised most frequent sense detection using word embeddings. In *DENVER*. Citeseer.

Sudha Bhingardive, Dharendra Singh, Rudra-murthy V, Hanumant Harichandra Redkar, and Pushpak Bhattacharyya. 2015b. Unsupervised most frequent sense detection using word embeddings. In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, May 31 - June 5, 2015*, pages 1238–1243.

Stephen Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python*. O’Reilly. (www.nltk.org/book).

Lars Borin, Markus Forsberg, and Lennart Lönngrén. 2013. Saldo: a touch of yin to WordNet’s yang. *Language Resources and Evaluation*, 47(4):1191–1211.

Fišer Darja, Jernej Novak, and Tomaž. 2012. sloWNet 3.0: development, extension and cleaning. In *Proceedings of 6th International Global Wordnet Conference (GWC 2012)*, pages 113–117. The Global WordNet Association.

Valéria de Paiva, Alexandre Rademaker, and Gerard de Melo. 2012. OpenWordNet-PT: an open Brazilian Wordnet for reasoning. EMap technical report, Escola de Matemática Aplicada, FGV, Brazil.

Sabri Elkateb, William Black, Horacio Rodríguez, Musa Alkhalifa, Piek Vossen, Adam Pease, and Christiane Fellbaum. 2006. Building a wordnet for Arabic. In *Proceedings of The fifth international conference on Language Resources and Evaluation (LREC 2006)*.

Christine Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database*. MIT Press.

Radovan Garabík and Indrė Pileckytė. 2013. From multilingual dictionary to lithuanian wordnet. In Katarína Gajdošová and Adriána Žáková, editors, *Natural Language Processing, Corpus Linguistics, E-Learning*, pages 74–80. Lüdenscheid: RAM-Verlag. http://korpus.juls.savba.sk/attachments/publications/lithuanian_wordnet_2013.pdf.

- Aitor Gonzalez-Agirre, Egoitz Laparra, and German Rigau. 2012. Multilingual central repository version 3.0: upgrading a very large lexical knowledge base. In *Proceedings of the 6th Global WordNet Conference (GWC 2012)*, Matsue.
- Donna Harman. 2011. Information retrieval evaluation. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, 3(2):1–119.
- Chu-Ren Huang, Shu-Kai Hsieh, Jia-Fei Hong, Yun-Zhu Chen, I-Li Su, Yong-Xiang Chen, and Sheng-Wei Huang. 2010. Chinese wordnet: Design and implementation of a cross-lingual knowledge processing infrastructure. *Journal of Chinese Information Processing*, 24(2):14–23. (in Chinese).
- Hitoshi Isahara, Francis Bond, Kiyotaka Uchimoto, Masao Utiyama, and Kyoko Kanzaki. 2008. Development of the Japanese WordNet. In *Sixth International conference on Language Resources and Evaluation (LREC 2008)*, Marrakech.
- Hong Jin Kang, Tao Chen, Muthu Kumar Chandrasekaran, and Min-Yen Kan. 2016. A comparison of word embeddings for english and cross-lingual chinese word sense disambiguation. *arXiv preprint arXiv:1611.02956*.
- Ravi Kumar and Sergei Vassilvitskii. 2010. Generalized distances between rankings. In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, pages 571–580, New York, NY, USA. ACM.
- Aiden Si Hong Lim. 2014. *Acquiring Predominant Word Senses in Multiple Languages*. Ph.D. thesis, School of Humanities and Social Sciences, Nanyang Technological University.
- Krister Lindén and Lauri Carlson. 2010. Finnwordnet — wordnet påfinska via översättning. *LexicoNordica — Nordic Journal of Lexicography*, 17:119–140. In Swedish with an English abstract.
- Quan Liu, Hui Jiang, Si Wei, Zhen-Hua Ling, and Yu Hu. 2015. Learning semantic word embeddings based on ordinal knowledge constraints. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, pages 1501–1511.
- Teng Long, Ryan Lowe, Jackie Chi Kit Cheung, and Doina Precup. 2016. Leveraging lexical resources for learning entity embeddings in multi-relational data. *arXiv preprint arXiv:1605.05416*.
- Diana McCarthy and John Carroll. 2003. Disambiguating nouns, verbs and adjectives using automatically acquired selectional preferences. *Computational Linguistics*, 29(4):639–654.
- Nurhil Hirfana Mohamed Noor, Suerya Sapuan, and Francis Bond. 2011. Creating the open Wordnet Bahasa. In *Proceedings of the 25th Pacific Asia Conference on Language, Information and Computation (PACLIC 25)*, pages 258–267, Singapore.
- Mortaza Montazery and Hesham Faili. 2010. Automatic Persian wordnet construction. In *23rd International conference on computational linguistics*, pages 846–850.
- Roberto Navigli and Simone Paolo Ponzetto. 2010. Babelnet: Building a very large multilingual semantic network. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pages 216–225.
- Roberto Navigli. 2009. Word sense disambiguation: A survey. *ACM Computing Surveys (CSUR)*, 41(2):10.
- Antoni Oliver, K. Šojat, and M. Srebačić. 2015. Automatic expansion of Croatian wordnet. In *Proceedings of the 29th CALS international conference “Applied Linguistic Research and Methodology”*, Zadar (Croatia).
- Noam Ordan and Shuly Wintner. 2007. Hebrew wordnet: a test case of aligning lexical databases across languages. *International Journal of Translation*, 19(1):39–58.
- Hamid Palangi, Li Deng, Yelong Shen, Jianfeng Gao, Xiaodong He, Jianshu Chen, Xinying Song, and Rabab Ward. 2016. Deep sentence embedding using long short-term memory networks: Analysis and application to information retrieval. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, 24(4):694–707.
- Alexander Panchenko. 2016. Best of both worlds: Making word sense embeddings interpretable. In *the 10th edition of the Language Resources and Evaluation Conference (LREC 2016)*.
- Ted Pedersen, Siddharth Patwardhan, and Jason Michelizzi. 2004. Wordnet:: Similarity: measuring the relatedness of concepts. In *Demonstration papers at HLT-NAACL 2004*, pages 38–41. Association for Computational Linguistics.
- Bolette Sandford Pedersen, Sanni Nimb, Jørg Asmussen, Nicolai Hartvig Sørensen, Lars Trap-Jensen, and Henrik Lorentzen. 2009. DanNet — the challenge of compiling a wordnet for Danish by reusing a monolingual dictionary. *Language Resources and Evaluation*, 43(3):269–299.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–1543.

- Emanuele Pianta, Luisa Bentivogli, and Christian Girardi. 2002. Multiwordnet: Developing an aligned multilingual database. In *Proceedings of the First International Conference on Global WordNet*, pages 293–302, Mysore, India.
- Maciej Piasecki, Stan Szpakowicz, and Bartosz Broda. 2009. *A Wordnet from the Ground Up*. Wrocław University of Technology Press. (ISBN 978-83-7493-476-3).
- Elisabete Pociello, Eneko Agirre, and Izaskun Aldezabal. 2011. Methodology and construction of the Basque wordnet. *Language Resources and Evaluation*, 45(2):121–142.
- Joel Pocostales. 2016. Nuig-unlp at semeval-2016 task 13: A simple word embedding-based approach for taxonomy extraction. *Proceedings of SemEval*, pages 1298–1302.
- Ida Raffaelli, Božo Bekavac, Željko Agić, and Marko Tadić. 2008. Building Croatian wordnet. In Attila Tanács, Dóra Csendes, Veronika Vincze, Christianne Fellbaum, and Piek Vossen, editors, *Proceedings of the Fourth Global WordNet Conference 2008*, pages 349–359, Szeged.
- Xin Rong. 2014. word2vec parameter learning explained. *arXiv preprint arXiv:1411.2738*.
- Sascha Rothe and Hinrich Schütze. 2015. Autoextend: Extending word embeddings to embeddings for synsets and lexemes. *arXiv preprint arXiv:1507.01127*.
- Ervin Ruci. 2008. On the current state of Albanet and related applications. Technical report, University of Vlora. (<http://fjalnet.com/technicalreportalbanet.pdf>).
- Benoît Sagot and Darja Fišer. 2008. Building a free French wordnet from multilingual resources. In European Language Resources Association (ELRA), editor, *Proceedings of the Sixth International Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco.
- Kiril Simov and Petya Osenova. 2010. Constructing of an ontology-based lexicon for Bulgarian. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, Mike Rosner, and Daniel Tapias, editors, *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta, may. European Language Resources Association (ELRA).
- Sofia Stamou, Goran Nenadic, and Dimitris Christodoulakis. 2004. Exploring Balkanet shared ontology for multilingual conceptual indexing. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC 2004)*, page 781–784, Lisbon.
- Liling Tan and Francis Bond. 2013. XLING: Matching query sentence to parallel corpus using topic models for word sense disambiguation. In *International Workshop on Semantic Evaluation (SemEval 2013)*.
- Duyu Tang, Furu Wei, Nan Yang, Ming Zhou, Ting Liu, and Bing Qin. 2014. Learning sentiment-specific word embedding for twitter sentiment classification. In *ACL (1)*, pages 1555–1565.
- Sareewan Thoongsup, Thatsanee Charoenporn, Kergrit Robkop, Tan Sinthurahat, Chumpol Mokrat, Virach Sornlertlamvanich, and Hitoshi Isahara. 2009. Thai wordnet construction. In *Proceedings of The 7th Workshop on Asian Language Resources (ALR7), Joint conference of the 47th Annual Meeting of the Association for Computational Linguistics (ACL) and the 4th International Joint Conference on Natural Language Processing (IJCNLP)*, Suntec, Singapore.
- Antonio Toral, Stefania Bracal, Monica Monachini, and Claudia Soria. 2010. Rejuvenating the Italian wordnet: upgrading, standardising, extending. In *Proceedings of the 5th International Conference of the Global WordNet Association (GWC-2010)*.
- Dan Tufiş, Radu Ion, Luigi Bozianu, Alexandru Ceaşu, and Dan Ştefănescu. 2008. Romanian wordnet: Current state, new applications and prospects. In *Proceedings of the 4th Global WordNet Association Conference*, pages 441–452, Szeged.
- Piek Vossen and Marten Postma. 2014. Open Dutch wordnet. In *Proceedings of the 7th Global WordNet Conference (GWC 2014)*, Tartu. (presentation only).
- Shan Wang and Francis Bond. 2013. Building the Chinese Open Wordnet (COW): Starting from core synsets. In *Proceedings of the 11th Workshop on Asian Language Resources, a Workshop at IJCNLP-2013*, pages 10–18, Nagoya.
- William Webber, Alistair Moffat, and Justin Zobel. 2010. A similarity measure for indefinite rankings. *ACM Trans. Inf. Syst.*, 28(4):20:1–20:38, nov.