

Semi-automatic WordNet Linking using Word Embeddings

Kevin Patel[†], Diptesh Kanojia^{†,♣,*}, Pushpak Bhattacharyya[†]

[†]Indian Institute of Technology Bombay, India

[♣]IITB-Monash Research Academy, India

^{*}Monash University, Australia

[†]{kevin.patel, diptesh, pb}@cse.iitb.ac.in

Abstract

Wordnets are rich lexico-semantic resources. Linked wordnets are extensions of wordnets, which link similar concepts in wordnets of different languages. Such resources are extremely useful in many Natural Language Processing (NLP) applications, primarily those based on knowledge-based approaches. In such approaches, these resources are considered as gold standard/oracle. Thus, it is crucial that these resources hold correct information. Thereby, they are created by human experts. However, manual maintenance of such resources is a tedious and costly affair. Thus techniques that can aid the experts are desirable. In this paper, we propose an approach to link wordnets. Given a synset of the source language, the approach returns a ranked list of potential candidate synsets in the target language from which the human expert can choose the correct one(s). Our technique is able to retrieve a winner synset in the top 10 ranked list for 60% of all synsets and 70% of noun synsets.

1 Introduction

Wordnets (Fellbaum, 1998) have been useful in different Natural Language Processing applications such as Word Sense Disambiguation (Tufiş et al., 2004; Sinha et al., 2006), Machine Translation (Knight and Luk, 1994) etc.

Linked Wordnets are extensions of wordnets. In addition to language specific information captured in constituent wordnets, linked wordnets have a notion of an interlingual index, which connects similar concepts in different languages. Such linked wordnets have found their application in machine translation (Hovy, 1998), cross-lingual information retrieval (Gonzalo et al., 1998), etc.

Given the extensive application of wordnets in different NLP applications, maintenance of wordnets involves expert involvement. Such involvement is costly both in terms of time and resources. This is further amplified in case of linked wordnets, where experts need to have knowledge of multiple languages. Thus, techniques that can help reduce the effort needed by experts are desirable.

Recently, deep learning has been extremely successful in a wide array of NLP applications. This is primarily due to the development of word embeddings, which have become a crucial component in modern NLP. They are learned in an unsupervised manner from large amounts of raw corpora. Bengio et al. (2003) were the first to propose neural word embeddings. Many word embedding models have been proposed since then (Collobert and Weston, 2008; Huang et al., 2012; Mikolov et al., 2013c; Levy and Goldberg, 2014). They have been efficiently utilized in many NLP applications: Part of Speech Tagging (Collobert and Weston, 2008), Named Entity Recognition (Collobert and Weston, 2008), Sentence Classification (Kim, 2014), Sentiment Analysis (Liu et al., 2015), Sarcasm Detection (Joshi et al., 2016)

Mikolov et al. (2013a) made a particularly interesting observation about the structure of the embedding space of different languages. They noted that there is a linear mapping between such spaces.

In this paper, we address the following question:

“Can information about the structure of embedding spaces of different languages and the relation among them be used to aid linking of corresponding wordnets?”

We demonstrate that this is true at least in the case of English and Hindi WordNets. We propose an approach to link them using word embeddings. Given a synset of the source language, the approach provides a ranked list of target synsets. This makes the overall linking task easy for human




Figure 1: Word embeddings of numbers and animals in English (left) and Spanish (right) (taken from (Mikolov et al., 2013a)).

experts, as they have to choose from a relatively small set of potential candidates. Our evaluation shows that our technique is able to retrieve a winner synset in the top 10 ranked list for 60% and 70% of all synsets and noun synsets respectively.

2 Background and Related Work

Princeton WordNet or the English WordNet was the first wordnet and inspired the development of many other wordnets. EuroWordNet (Vossen and others, 1997) is a linked wordnet comprising of wordnets for European languages, *viz.*, Dutch, Italian, Spanish, German, French, Czech and Estonian. Each of these wordnets is structured in the same way as the Princeton WordNet for English (Miller et al., 1990) - synsets (sets of synonymous words) and semantic relations between them. Each wordnet separately captures a language-specific information. In addition, the wordnets are linked to an Inter-Lingual-Index, which uses Princeton WordNet as a base. This index enables one to go from concepts in one language to similar concepts in any other language. Such features make this resource helpful in cross-lingual NLP applications.

IndoWordNet (Bhattacharyya, 2010) is a linked wordnet comprising of wordnets for major In-

dian languages, *viz.*, Assamese, Bengali, Bodo, Gujarati, Hindi, Kannada, Kashmiri, Konkani, Malayalam, Manipuri, Marathi, Nepali, Oriya, Punjabi, Sanskrit, Tamil, Telugu, and Urdu. These wordnets have been created using the expansion approach using Hindi WordNet as a pivot, which is partially linked to English WordNet. Previously, Joshi et al. (2012a) come up with a heuristic based measure where they use bilingual dictionaries to link two wordnets. They combine scores using various heuristics and generate a list of potential candidates for linked synsets.

Singh et al. (2016) discuss a method to improve the current status of Hindi-English linkage and present a generic methodology *i.e.*, manually creating bilingual mappings for concepts which are unavailable in either of the languages or not present as a synset in the target wordnet. Their method is beneficial for culture-specific synsets, or for non-existing concepts; but, it is cost and time inefficient, and requires a lot of manual effort on the part of a lexicographer.

Our approach is mainly geared towards reducing effort on the part of the lexicographers.

3 Problem Statement

Given wordnets of two different languages E and F with sets of synsets $\{s_E^1, s_E^2, \dots, s_E^m\}$ and $\{s_F^1, s_F^2, \dots, s_F^n\}$ respectively, find mappings of the form $\langle s_E^i, s_F^j \rangle$ which are semantically correct.

4 Approach

We adapted the technique of translating words in Mikolov et al. (2013a) to translate synsets (see fig 1). In order to do so, however, we need "synset embeddings". We computed the same by assigning to a synset-id, the average of the "word embeddings" of its synset-members. To the best of our knowledge, this is a first attempt at solving this problem using word embeddings. The following is a detailed description of the technique.

Let E and F be two languages. Let $|E|$ and $|F|$ be the number of synsets in wordnets of E and F respectively. Let s_E^i and s_F^j be the i^{th} and j^{th} synsets of E and F respectively, with $s_E^i = \{e_\alpha^1, e_\alpha^2, \dots, e_\alpha^{m_i}\}$ and $s_F^j = \{f_\beta^1, f_\beta^2, \dots, f_\beta^{n_j}\}$, where e_α^p and f_β^q are words in vocabulary of E and F respectively for $1 \leq p \leq m_i$ and $1 \leq q \leq n_j$, and $1 \leq i \leq |E|$ and $1 \leq j \leq |F|$.

Let $v_{e_\alpha^p}$ be the word embedding corresponding to e_α^p . Then we estimate $v_{s_E^i}$, the embedding for synset s_E^i , as

$$v_{s_E^i} = \frac{1}{m_i} \sum_{p=0}^{m_i} v_{e_\alpha^p} \quad (1)$$

Similarly,

$$v_{s_F^j} = \frac{1}{n_j} \sum_{q=0}^{n_j} v_{f_\beta^q} \quad (2)$$

Given links of the form $\langle s_E^i, s_F^j \rangle$, we learn W such that the error Err

$$Err = \|W.v_{s_E^i} - v_{s_F^j}\|^2 \quad (3)$$

is minimized.

Now, to find a mapping for a new synset s_E^k , one needs to

1. Calculate $v' = W.v_{s_E^k}$
2. Find $v_{s_F^l}$ such that $v_{s_F^l} \cdot v'$ is maximized
3. Create link $\langle s_E^k, s_F^l \rangle$

Our hypothesis is that for a given synset-id, the noise added to its representative embedding by a highly polysemous synset-member will be canceled out, while the actual information content pertaining to that synset-id will be enhanced, due to contribution from other, relatively less polysemous, synset members.

5 Experiments

Datasets

We applied our technique to link Hindi and English Wordnets. We obtained a dataset of mappings between English and Hindi wordnets from the developers of IndoWordNet. These mappings are of the form $\langle hindi_synset_id, english_synset_id, link_type \rangle$, where $link_type \in \{\text{DIRECT}, \text{HYPERNYMY}, \text{etc.}\}$. For this experiment, we focused solely on DIRECT links. There are a total of 6,883 such mappings, the distribution among classes of which is mentioned in table 1

Class	Count
Noun	4757
Adjective	1283
Verb	680
Adverb	143

Table 1: Distribution of available links among various classes

For the English language, we used the pre-trained word embeddings published by Google that were trained on part of Google News Dataset (about 100 billion tokens). These embeddings are of dimension 300, and are created using CBOW model with negative sampling. For the Hindi language, we trained word embeddings on BOJAR HindMonoCorp dataset (Bojar et al., 2014). Mikolov et al. (2013b) suggests that the input embeddings' dimension should be at least 2.5 to 4 times that of the output dimension. But we also wanted to check what happens when they are equal. Therefore, we trained two sets of embeddings, one of dimension 300, and the other of dimension 1200.

Evaluation Metric

We use the accuracy@ n measure, i.e the prediction is said to be correct if one out of the top n results returned is correct. This is because accuracy@1 is an underestimate of the system's per-

formance, as higher-ranking synonym translations will be counted as mistakes.

	Predicted Label	Accuracy @1	Accuracy @3	Accuracy @5
True label	Prediction1			
	Prediction2			
	Prediction3			
	Prediction4			
	Prediction5			

Figure 2: Accuracy@n: The green colored cells indicate the predictions considered for exact match for a given accuracy@n

6 Results and Discussion

Table 2 shows the overall accuracy@n of the system, for different values of n. We also performed a per word-class evaluation, along with different settings for the embedding dimensions. Table 3 and Table 4 shows the accuracy for different word classes ¹.

Acc@1	Acc@3	Acc@5	Acc@8	Acc@10
0.29	0.45	0.52	0.58	0.60

Table 2: Results for the overall setting: Dimension of English embeddings=300, Dimensions of Hindi embeddings=300

Word Class	Acc@1	Acc@3	Acc@5	Acc@8	Acc@10
Noun	0.35	0.53	0.60	0.65	0.67
Adjective	0.26	0.44	0.50	0.57	0.60
Verb	0.15	0.25	0.29	0.33	0.37
Adverb	0.28	0.51	0.59	0.70	0.73

Table 3: Results for the setting: Dimension of English embeddings=300, Dimensions of Hindi embeddings=300

We observe that except for verbs, the approach performs decently. Here we mention some of the reasons for poor performance, as well as possible methods to address them.

- The approach to create synset embeddings is inadequate. The current averaging approach only takes the synset members into account, while ignoring gloss and examples, which could provide additional information. A potential candidate approach for creating synset embeddings should properly utilize the set of

¹All values reported are the average values obtained from 3-fold cross validation.

Word Class	Acc@1	Acc@3	Acc@5	Acc@8	Acc@10
Noun	0.35	0.52	0.58	0.63	0.66
Adjective	0.12	0.20	0.24	0.30	0.32
Verb	0.17	0.27	0.32	0.35	0.39
Adverb	0.38	0.52	0.65	0.76	0.80

Table 4: Results for the setting: Dimension of English embeddings=300, Dimensions of Hindi embeddings=1200

French synonyms, gloss, example sentences, and synset relations.

- Synset members are often phrases instead of words. Creating phrase embeddings is a different problem altogether.
- Currently, we utilized a word embedding model which gives only one embedding per word. That is one of the reasons for ambiguity. A model which provides one embedding per sense of a word will be a more appropriate.
- The linear transformation approach is incorrect. While (Mikolov et al., 2013a) shows the linear relation between English and Spanish languages, this may not be true for all pairs of languages.
- Perhaps, something is fundamentally missing in word embeddings. Probably presence of only co-occurrence information and lack of other information such as word order, argument frames(for verbs), *etc.* leads to this poor performance.

However, we were unable to find an explanation for the degradation of results of adjectives when using 1200 dimensions for Hindi word embeddings.

7 Conclusion and Future Work

In this paper, we described an approach to link wordnets. It entails creating synset embeddings using the word embeddings of the synset members, and learning a function to map the embedding of a synset from the source language to an embedding in the space of target language, and returning the nearest neighbors as potential candidates for linking. Our evaluation shows that our technique is able to retrieve a winner synset in the top 10 ranked list for 60% and 70% of all synsets and noun synsets, respectively. Although, it did

not achieve significantly good results for other classes, especially verbs. We discussed the possible reasons for poor performance and suggested mechanisms to address the same.

In future, we plan to continue this work, and explore each of the above possible reasons for poor performance, in order to mitigate them. We will also evaluate it in an active learning setting. Eventually, we aim to integrate our work with tools such as the ones created by Joshi et al. (2012b), *etc.* so that our work can be used by lexicographers and researchers alike.

References

- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. A neural probabilistic language model. *J. Mach. Learn. Res.*, 3:1137–1155, March.
- Pushpak Bhattacharyya. 2010. Indowordnet. In *In Proc. of LREC-10*. Citeseer.
- Ondřej Bojar, Vojtěch Diatka, Pavel Rychlý, Pavel Straňák, Vít Suchomel, Aleš Tamchyna, and Daniel Zeman. 2014. HindMonoCorp 0.5.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: deep neural networks with multitask learning. In William W. Cohen, Andrew McCallum, and Sam T. Roweis, editors, *ICML*, volume 307 of *ACM International Conference Proceeding Series*, pages 160–167. ACM.
- Christiane Fellbaum. 1998. *WordNet*. Wiley Online Library.
- Julio Gonzalo, Felisa Verdejo, Irina Chugur, and Juan Cigarran. 1998. Indexing with wordnet synsets can improve text retrieval. *arXiv preprint cmp-lg/9808002*.
- Eduard Hovy. 1998. Combining and standardizing large-scale, practical ontologies for machine translation and other uses. In *Proceedings of the 1st International Conference on Language Resources and Evaluation (LREC)*, pages 535–542.
- Eric H. Huang, Richard Socher, Christopher D. Manning, and Andrew Y. Ng. 2012. Improving Word Representations via Global Context and Multiple Word Prototypes. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Salil Joshi, Arindam Chatterjee, Arun Karthikeyan Karra, and Pushpak Bhattacharyya. 2012a. Eating your own cooking: automatically linking wordnet synsets of two languages.
- Salil Joshi, Arindam Chatterjee, Karthikeyan Arun Karra, and Pushpak Bhattacharyya. 2012b. Eating your own cooking: Automatically linking wordnet synsets of two languages. In *Proceedings of COLING 2012: Demonstration Papers*, pages 239–246. The COLING 2012 Organizing Committee.
- Aditya Joshi, Vaibhav Tripathi, Kevin Patel, Pushpak Bhattacharyya, and Mark Carman. 2016. Are word embedding-based features useful for sarcasm detection? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1006–1011, Austin, Texas, November. Association for Computational Linguistics.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar, October. Association for Computational Linguistics.
- Kevin Knight and Steve K Luk. 1994. Building a large-scale knowledge base for machine translation. In *AAAI*, volume 94, pages 773–778.
- Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 2: Short Papers*, pages 302–308.
- Pengfei Liu, Shafiq R Joty, and Helen M Meng. 2015. Fine-grained opinion mining with recurrent neural networks and word embeddings. In *EMNLP*, pages 1433–1443.
- Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. 2013a. Exploiting similarities among languages for machine translation. *CoRR*, abs/1309.4168.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013b. Distributed representations of words and phrases and their compositionality. *CoRR*, abs/1310.4546.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013c. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- George A Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine J Miller. 1990. Introduction to wordnet: An on-line lexical database. *International journal of lexicography*, 3(4):235–244.
- Meghna Singh, Rajita Shukla, Jaya Jha, Laxmi Kashyap, Diptesh Kanojia, and Pushpak Bhattacharyya. 2016. Mapping it differently: A solution to the linking challenges. In *Eighth Global Wordnet Conference*. GWC 2016.
- Manish Sinha, Mahesh Reddy, and Pushpak Bhattacharyya. 2006. An approach towards construction and application of multilingual indo-wordnet. In *3rd*

Global Wordnet Conference (GWC 06), Jeju Island, Korea.

Dan Tufiș, Radu Ion, and Nancy Ide. 2004. Fine-grained word sense disambiguation based on parallel corpora, word alignment, word clustering and aligned wordnets. In *Proceedings of the 20th international conference on Computational Linguistics*, page 1312. Association for Computational Linguistics.

Piek Vossen et al. 1997. Eurowordnet: a multilingual database for information retrieval. In *Proceedings of the DELOS workshop on Cross-language Information Retrieval*, pages 5–7.