

# Enhanced Word Representations for Bridging Anaphora Resolution

Yufang Hou  
IBM Research Ireland  
yhou@ie.ibm.com

## Abstract

Most current models of word representations (e.g., *GloVe*) have successfully captured fine-grained semantics. However, semantic similarity exhibited in these word embeddings is not suitable for resolving bridging anaphora, which requires the knowledge of associative similarity (i.e., relatedness) instead of semantic similarity information between synonyms or hypernyms. We create word embeddings (*embeddings\_PP*) to capture such relatedness by exploring the syntactic structure of noun phrases. We demonstrate that using *embeddings\_PP* alone achieves around 30% of accuracy for bridging anaphora resolution on the ISNotes corpus. Furthermore, we achieve a substantial gain over the state-of-the-art system (Hou et al., 2013b) for bridging antecedent selection.

## 1 Introduction

Bridging (Clark, 1975; Prince, 1981; Gundel et al., 1993) establishes entity coherence in a text by linking anaphors and antecedents via various non-identity relations. In Example 1, the link between the bridging anaphor (**the chief cabinet secretary**) and the antecedent (*Japan*) establish local (entity) coherence.

(1) Yet another political scandal is racking *Japan*. On Friday, **the chief cabinet secretary** announced that **eight cabinet ministers** had received five million yen from the industry.

Choosing the right antecedents for bridging anaphors is a subtask of bridging resolution. For this subtask, most previous work (Poesio et al., 2004; Lassalle and Denis, 2011; Hou et al., 2013b) calculate semantic relatedness between an anaphor and its antecedent based on word co-occurrence count using certain syntactic patterns.

Most recently, word embeddings gain a lot popularity in NLP community because they reflect human intuitions about semantic similarity and relatedness. Most word representation models explore the distributional hypothesis which states that words occurring in similar contexts have similar meanings (Harris, 1954). State-of-the-art word representations such as word2vec skipgram (Mikolov et al., 2013) and GloVe (Pennington et al., 2014) have been shown to perform well across a variety of NLP tasks, including textual entailment (Rocktäschel et al., 2016), reading comprehension (Chen et al., 2016), and information status classification (Hou, 2016). However, these word embeddings capture both “genuine” similarity and relatedness, and they may in some cases be detrimental to downstream performance (Kielbaso et al., 2015). Bridging anaphora resolution is one of such cases which requires lexical association knowledge instead of semantic similarity information between synonyms or hypernyms. In Example 1, among all antecedent candidates, “*the chief cabinet secretary*” is the most similar word to the bridging anaphor “**eight cabinet ministers**” but obviously it is not the antecedent for the latter.

In this paper, we explore the syntactic structure of noun phrases (NPs) to derive contexts for nouns in the GloVe model. We find that the prepositional structure (e.g., **X of Y**) and the possessive structure (e.g., **Y’s X**) are a useful context source for the representation of nouns in terms of relatedness for bridging relations.

We demonstrate that using our word embeddings based on PP contexts (*embeddings\_PP*) alone achieves around 30% of accuracy on bridging anaphora resolution in the ISNotes corpus, which is 12% better than the original GloVe word embeddings. Moreover, adding an additional feature based on *embeddings\_PP* leads to a significant improvement over a state-of-the-art system on

bridging anaphora resolution (Hou et al., 2013b).

## 2 Related Work

**Bridging anaphora resolution.** Anaphora plays an important role in discourse comprehension. Different from *identity anaphora* which indicates that a noun phrase refers back to the same entity introduced by previous descriptions in the discourse, *bridging anaphora* links anaphors and antecedents via lexico-semantic, frame or encyclopedic relations.

Bridging resolution has to recognize bridging anaphors and find links to antecedents. There has been a few works tackling full bridging resolution (Hahn et al., 1996; Hou et al., 2014). In recent years, various computational approaches have been developed for bridging anaphora recognition (Markert et al., 2012; Hou et al., 2013a) and for bridging antecedent selection (Poesio et al., 2004; Hou et al., 2013b). This work falls into the latter category and we create a new lexical knowledge resource for the task of choosing antecedents for bridging anaphors.

Previous work on bridging anaphora resolution (Poesio et al., 2004; Lassalle and Denis, 2011; Hou et al., 2013b) explore word co-occurrence count in certain syntactic preposition patterns to calculate word relatedness. These patterns encode associative relations between nouns which cover a variety of bridging relations. Our PP context model exploits the same principle but is more general. Unlike previous work which only consider a small number of prepositions per anaphor, the PP context model considers all prepositions for all nouns in big corpora. It also includes the possessive structure of NPs. The resulting word embeddings are a general resource for bridging anaphora resolution. In addition, it enables efficient computation of word association strength through low-dimensional matrix operations.

**Enhanced word embeddings.** Recently, a few approaches investigate different ways to improve the vanilla word embeddings. Levy and Goldberg (2014) explore the dependency-based contexts in the Skip-Gram model. The authors replace the linear bag-of-words contexts in the original Skip-Gram model with the syntactic contexts derived from the automatically parsed dependency trees. They observe that the dependency-based embeddings exhibit more functional similarity than the original skip-gram embeddings. Heinzerling et al.

(2017) show that incorporating dependency-based word embeddings into their selectional preference model slightly improve coreference resolution performance. Kiela et al. (2015) try to learn word embeddings for similarity and relatedness separately by utilizing a thesaurus and a collection of psychological association norms. The authors report that their relatedness-specialized embeddings perform better on document topic classification than similarity embeddings. Schwartz et al. (2016) demonstrate that symmetric patterns (e.g. X or Y) are the most useful contexts for the representation of verbs and adjectives. Our work follows in this vein and we are interested in learning word representations for bridging relations.

## 3 Approach

### 3.1 Asymmetric Prepositional and Possessive Structures

The syntactic prepositional and possessive structures of NPs encode a variety of bridging relations between anaphors and their antecedents. For instance, *the rear door of that red car* indicates the part-of relation between “door” and “car”, and *the company’s new appointed chairman* implies the employment relation between “chairman” and “company”. We therefore extract noun pairs *door-car*, *chairman-company* by using syntactic structure of NPs which contain prepositions or possessive forms.

It is worth noting that bridging relations expressed in the above syntactic structures are asymmetric. So for each noun pair, we keep the head on the left and the noun modifier on the right. However, a lot of nouns can appear on both positions, such as “*travelers in the train station*”, “*travelers from the airport*”, “*hotels for travelers*”, “*the destination for travelers*”. To capture the differences between these two positions, we add the postfix “\_PP” to the nouns on the left. Thus we extract the following four pairs from the above NPs: *travelers\_PP-station*, *travelers\_PP-airport*, *hotels\_PP-travelers*, *destination\_PP-travelers*.

### 3.2 Word Embeddings Based on PP Contexts (*embeddings\_PP*)

Our PP context model is based on GloVe (Pennington et al., 2014), which obtains state-of-the-art results on various NLP tasks. We extract noun pairs as described in Section 3.1 from the automatically parsed Gigaword corpus (Parker et al., 2011;

Target Word	<i>embeddings_PP</i>	<i>GloVe_Giga</i>
president	minister, mayor, governor, clinton bush	vice, presidency, met, former presidents
president_PP	vice-president_PP, federation, republic usa, corporation	—
residents	villagers, citizens, inhabitants, families participants	locals, villagers, people, citizens homes
residents_PP	resident_PP, neighborhood, shemona <sup>1</sup> ashraf, suburbs	—
members	participants, leaders, colleagues, officials lawmakers	member, representatives, others, leaders groups
members_PP	member_PP, representatives_PP, basij <sup>2</sup> leaders_PP, community	—
travelers	travellers, thirsts_PP, shoppers quarantines_PP, needle-sharing_PP	travellers, passengers, vacationers tourists, shoppers
travelers_PP	e-tickets, travellers_PP, cairngorms <sup>3</sup> flagstaffs_PP, haneda <sup>4</sup>	—

<sup>1</sup> Shemona is a city in Israel.    <sup>2</sup> Basij is a paramilitary group in Iran.

<sup>3</sup> Cairngorms is mountain range in Scotland.    <sup>4</sup> Haneda is an airport in Japan.

Table 1: Target words and their top five nearest neighbors in *embeddings\_PP* and *GloVe\_Giga*

Napoles et al., 2012). We treat each noun pair as a sentence containing only two words and concatenate all 197 million noun pairs in one document. We employ the GloVe toolkit<sup>1</sup> to train the PP context model on the above extracted noun pairs. All tokens are converted to lowercase, and words that appear less than 10 times are filtered. This results in a vocabulary of around 276k words and 188k distinct nouns without the postfix “\_PP”. We set the context window size as two and keep other parameters the same as in Pennington et al. (2014). We report results for 100 dimension embeddings, though similar trends were also observed with 200 and 300 dimensions.

For comparison, we also trained a 100 dimension word embeddings (*GloVe\_Giga*) on the whole Gigaword corpus, using the same parameters reported in Pennington et al. (2014).

Table 1 lists a few target words and their top five nearest neighbors (using cosine similarity) in *embeddings\_PP* and *GloVe\_Giga* respectively. For the target words “residents” and “members”, both *embeddings\_PP* and *GloVe\_Giga* yield a list of similar words and most of them have the same semantic type as the target word. For the “travelers” example, *GloVe\_Giga* still presents the similar words with the same semantic type, while *embed-*

*dings\_PP* generates both similar words and related words (words containing the postfix “\_PP”). More importantly, it seems that *embeddings\_PP* can find reasonable semantic roles for nominal predicates (target words containing the postfix “\_PP”). For instance, “president\_PP” is mostly related to countries or organizations, and “residents\_PP” is mostly related to places.

The above examples can be seen as qualitative evaluation for our PP context model. We assume that *embeddings\_PP* can be served as a lexical knowledge resource for bridging antecedent selection. In the next section, we will demonstrate the effectiveness of *embeddings\_PP* for the task of bridging anaphora resolution.

## 4 Quantitative Evaluation

For the task of bridging anaphora resolution, we use the dataset ISNotes<sup>2</sup> released by Hou et al. (2013b). This dataset contains around 11,000 NPs annotated for information status including 663 bridging NPs and their antecedents in 50 texts taken from the WSJ portion of the OntoNotes corpus (Weischedel et al., 2011). It is notable that bridging anaphors in ISNotes are not limited to definite NPs as in previous work (Poesio et al., 1997, 2004; Lassalle and Denis, 2011).

<sup>1</sup><https://github.com/stanfordnlp/GloVe>

<sup>2</sup><http://www.h-its.org/english/research/nlp/download>

The semantic relations between anaphor and antecedent in the corpus are quite diverse: only 14% of anaphors have a part-of/attribute-of relation with the antecedent and only 7% of anaphors stand in a set relationship to the antecedent. 79% of anaphors have “other” relation with their antecedents, without further distinction. This includes encyclopedic relations such as **the waiter** – *restaurant* as well as context-specific relations such as **the thieves** – *palms*.

We follow Hou et al. (2013b)’s experimental setup and reimplement *MLN model II* as our baseline. We first test the effectiveness of *embeddings\_PP* alone to resolve bridging anaphors. Then we show that incorporating *embeddings\_PP* into *MLN model II* significantly improves the result.

#### 4.1 Using *embeddings\_PP* Alone

For each anaphor  $a$ , we simply construct the list of antecedent candidates  $E_a$  using NPs preceding  $a$  from the same sentence as well as from the previous two sentences. Hou et al. (2013b) found that globally salient entities are likely to be the antecedents of all anaphors in a text. We approximate this by adding NPs from the first sentence of the text to  $E_a$ . This is motivated by the fact that ISNotes is a newswire corpus and globally salient entities are often introduced in the beginning of an article. On average, each bridging anaphor has 19 antecedent candidates using this simple antecedent candidate selection strategy.

Given an anaphor  $a$  and its antecedent candidate list  $E_a$ , we predict the most related NP among all NPs in  $E_a$  as the antecedent for  $a$ . The relatedness is measured via cosine similarity between the head of the anaphor (plus the postfix “\_PP”) and the head of the candidate.

This simple deterministic approach based on *embeddings\_PP* achieves an accuracy of 30.32% on the ISNotes corpus. Following Hou et al. (2013b), accuracy is calculated as the proportion of the correctly resolved bridging anaphors out of all bridging anaphors in the corpus.

We found that using *embeddings\_PP* outperforms using other word embeddings by a large margin (see Table 2), including the original GloVe vectors trained on Gigaword and Wikipedia 2014 dump (*GloVe\_GigaWiki14*) and GloVe vectors that we trained on Gigaword only (*GloVe\_Giga*). This confirms our observation in Section 3.2 that *em-*

	acc
<i>GloVe_GigaWiki14</i>	18.10
<i>GloVe_Giga</i>	19.00
<i>embeddings_wo_PPSuffix</i>	22.17
<i>embeddings_PP</i>	<b>30.32</b>

Table 2: Results of *embeddings\_PP* alone for bridging anaphora resolution compared to the baselines. Bold indicates statistically significant differences over the baselines using randomization test ( $p < 0.01$ ).

*biddings\_PP* can capture the relatedness between anaphor and antecedent for various bridging relations.

To understand the role of the suffix “\_PP” in *embeddings\_PP*, we trained word vectors *embeddings\_wo\_PPSuffix* using the same noun pairs as in *embeddings\_PP*. For each noun pair, we remove the suffix “\_PP” attached to the head noun. We found that using *embeddings\_wo\_PPSuffix* only achieves an accuracy of 22.17% (see Table 2). This indicates that the suffix “\_PP” is the most significant factor in *embeddings\_PP*. Note that when calculating cosine similarity based on the first three word embeddings in Table 2, we do not add the suffix “\_PP” to the head of an bridging anaphor because such words do not exist in these word vectors.

#### 4.2 *MLN model II + embeddings\_PP*

*MLN model II* is a joint inference framework based on Markov logic networks (Domingos and Lowd, 2009). In addition to modeling the semantic, syntactic and lexical constraints between the anaphor and the antecedent (local constraints), it models that:

- semantically or syntactically related anaphors are likely to share the same antecedent (joint inference constraints);
- a globally salient entity is preferred to be the antecedent of all anaphors in a text even if the entity is distant to the anaphors (global salience constraints);
- several bridging relations are strongly signaled by the semantic classes of the anaphor and the antecedent, e.g., a job title anaphor such as *chairman* prefers a GPE or an organization antecedent (semantic class constraints).

	<b>acc</b>
<i>MLN model II</i>	41.32
<i>MLN model II + GloVe_GigaWiki14</i>	39.52
<i>MLN model II + embeddings_wo_PPSuffix</i>	40.42
<i>MLN model II + embeddings_PP</i>	<b>45.85</b>

Table 3: Results of integrating *embeddings\_PP* into *MLN model II* for bridging anaphora resolution compared to the baselines. Bold indicates statistically significant differences over the baselines using randomization test ( $p < 0.01$ ).

Due to the space limit, we omit the details of *MLN model II*, but refer the reader to Hou et al. (2013b) for a full description.

We add one constraint into *MLN model II* based on *embeddings\_PP*: each bridging anaphor  $a$  is linked to its most related antecedent candidate using cosine similarity. We use the same strategy as in the previous section to construct the list of antecedent candidates for each anaphor. Unlike the previous section, which only uses the vector of the NP head to calculate relatedness, here we include all common nouns occurring before the NP head as well because they also represent the core semantic of an NP (e.g., “*earthquake victims*” and “*the state senate*”).

Specifically, given an NP, we first construct a list  $N$  which consists of the head and all common nouns appearing before the head, we then represent the NP as a vector  $v$  using the following formula, where the suffix “\_PP” is added to each  $n$  if the NP is a bridging anaphor:

$$v = \frac{\sum_{n \in N} \text{embeddings\_PP}_n}{|N|} \quad (1)$$

Table 3 shows that adding the constraint based on *embeddings\_PP* improves the result of *MLN model II* by 4.5%. However, adding the constraint based on the vanilla word embeddings (*GloVe\_GigaWiki14*) or the word embeddings without the suffix “\_PP” (*embeddings\_wo\_PPSuffix*) slightly decreases the result compared to *MLN model II*. Although *MLN model II* already explores preposition patterns to calculate relatedness between head nouns of NPs, it seems that the feature based on *embeddings\_PP* is complementary to the original preposition pattern feature. Furthermore, the vector model allows us to represent the meaning of an NP beyond its head easily.

## 5 Conclusions

We present a PP context model based on GloVe by exploring the asymmetric prepositional structure (e.g.,  $\mathbf{X}$  of  $\mathbf{Y}$ ) and possessive structure (e.g.,  $\mathbf{Y}$ ’s  $\mathbf{X}$ ) of NPs. We demonstrate that the resulting word vectors (*embeddings\_PP*) are able to capture the relatedness between anaphor and antecedent in various bridging relations. In addition, adding the constraint based on *embeddings\_PP* yields a significant improvement over a state-of-the-art system on bridging anaphora resolution in ISNotes (Hou et al., 2013b).

For the task of bridging anaphora resolution, Hou et al. (2013b) pointed out that future work needs to explore wider context to resolve context-specific bridging relations. Here we combine the semantics of pre-nominal modifications and the head by vector average using *embeddings\_PP*. We hope that our embedding resource<sup>3</sup> will facilitate further research into improved context modeling for bridging relations.

## Acknowledgments

The author thanks the anonymous reviewers for their valuable feedback.

## References

- Danqi Chen, Jason Bolton, and Christopher D. Manning. 2016. A thorough examination of the CNN/Daily mail reading comprehension task. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Berlin, Germany, 7–12 August 2016. pages 2358–2367.
- Herbert H. Clark. 1975. Bridging. In *Proceedings of the Conference on Theoretical Issues in Natural Language Processing*, Cambridge, Mass., June 1975. pages 169–174.

<sup>3</sup>*embeddings\_PP* can be downloaded from <https://doi.org/10.5281/zenodo.1211616>

- Pedro Domingos and Daniel Lowd. 2009. *Markov Logic: An Interface Layer for Artificial Intelligence*. Morgan Claypool Publishers.
- Jeanette K. Gundel, Nancy Hedberg, and Ron Zacharski. 1993. Cognitive status and the form of referring expressions in discourse. *Language* 69:274–307.
- Udo Hahn, Michael Strube, and Katja Markert. 1996. **Bridging textual ellipses**. In *Proceedings of the 16th International Conference on Computational Linguistics*, Copenhagen, Denmark, 5–9 August 1996. volume 1, pages 496–501. <http://www.aclweb.org/anthology/C96-1084.pdf>.
- Zellig S. Harris. 1954. Distributional structure. *Word* 10:146–162.
- Benjamin Heinzerling, Nafise Sadat Moosavi, and Michael Strube. 2017. Revisiting selectional preferences for coreference resolution. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Copenhagen, Denmark, 7–11 November 2017. pages 1332–1339.
- Yufang Hou. 2016. Incremental fine-grained information status classification using attention-based LSTMs. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016. pages 1880–1890.
- Yufang Hou, Katja Markert, and Michael Strube. 2013a. **Cascading collective classification for bridging anaphora recognition using a rich linguistic feature set**. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013. pages 814–820. <http://aclweb.org/anthology/D13-1077.pdf>.
- Yufang Hou, Katja Markert, and Michael Strube. 2013b. **Global inference for bridging anaphora resolution**. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, 9–14 June 2013. pages 907–917. <http://aclweb.org/anthology/N13-1111.pdf>.
- Yufang Hou, Katja Markert, and Michael Strube. 2014. **A rule-based system for unrestricted bridging resolution: Recognizing bridging anaphora and finding links to antecedents**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, Doha, Qatar, 25–29 October 2014. pages 2082–2093. <http://aclweb.org/anthology/D13-1077.pdf>.
- Douwe Kiela, Felix Hill, and Stephen Clark. 2015. Specializing word embeddings for similarity or relatedness. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Lisbon, Portugal, 17–21 September 2015. pages 2044–2048.
- Emmanuel Lassalle and Pascal Denis. 2011. Leveraging different meronym discovery methods for bridging resolution in French. In *Proceedings of the 8th Discourse Anaphora and Anaphor Resolution Colloquium (DAARC 2011)*, Faro, Algarve, Portugal, 6–7 October 2011. pages 35–46.
- Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *Proceedings of the 52st Annual Meeting of the Association for Computational Linguistics*, Baltimore, USA, 22–27 June 2014.
- Katja Markert, Yufang Hou, and Michael Strube. 2012. **Collective classification for fine-grained information status**. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, Jeju Island, Korea, 8–14 July 2012. pages 795–804. <http://www.aclweb.org/anthology/P12-1084.pdf>.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26 (NIPS 2013)*, pages 3111–3119.
- Courtney Napoles, Matthew Gormley, and Benjamin Van Durme. 2012. Annotated Gigaword. In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction & Web-scale Knowledge Extraction (AKBC-WEKEX)* Montréal, Québec, Canada, 7-8 June 2012. pages 95–100.
- Robert Parker, David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2011. English Gigaword Fifth Edition. LDC2011T07.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, Doha, Qatar, 25–29 October 2014. pages 1532–1543.
- Massimo Poesio, Rahul Mehta, Axel Maroudas, and Janet Hitzeman. 2004. Learning to resolve bridging references. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*, Barcelona, Spain, 21–26 July 2004. pages 143–150.
- Massimo Poesio, Renata Vieira, and Simone Teufel. 1997. Resolving bridging references in unrestricted text. In *Proceedings of the ACL Workshop on Operational Factors in Practical, Robust Anaphora Resolution for Unrestricted Text*, Madrid, Spain, July 1997. pages 1–6.
- Ellen F. Prince. 1981. Towards a taxonomy of given-new information. In P. Cole, editor, *Radical Pragmatics*, Academic Press, New York, N.Y., pages 223–255.

Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomas Kocisky, and Phil Blunsom. 2016. Reasoning about entailment with neural attention. In *Proceedings of the 4th International Conference on Learning Representations*, San Juan, Puerto Rico, 2-4 May 2016.

Roy Schwartz, Roi Reichart, and Ari Rappoport. 2016. Symmetric patterns and coordinations: Fast and enhanced representations of Verbs and Adjectives. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, San Diego, California, 12–17 June 2016. pages 499–505.

Ralph Weischedel, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Nianwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, Mohammed El-Bachouti, Robert Belvin, and Ann Houston. 2011. OntoNotes release 4.0. LDC2011T03, Philadelphia, Penn.: Linguistic Data Consortium.