

A Deterministic Algorithm for Bridging Anaphora Resolution

Yufang Hou

IBM Research Ireland

yhou@ie.ibm.com

Abstract

Previous work on bridging anaphora resolution (Poesio et al., 2004; Hou et al., 2013b) use syntactic preposition patterns to calculate word relatedness. However, such patterns only consider NPs’ head nouns and hence do not fully capture the semantics of NPs. Recently, Hou (2018) created word embeddings (*embeddings_PP*) to capture associative similarity (i.e., relatedness) between nouns by exploring the syntactic structure of noun phrases. But *embeddings_PP* only contains word representations for nouns. In this paper, we create new word vectors by combining *embeddings_PP* with *GloVe*. This new word embeddings (*embeddings_bridging*) are a more general lexical knowledge resource for bridging and allow us to represent the meaning of an NP beyond its head easily. We therefore develop a deterministic approach for bridging anaphora resolution, which represents the semantics of an NP based on its head noun and modifications. We show that this simple approach achieves the competitive results compared to the best system in Hou et al. (2013b) which explores Markov Logic Networks to model the problem. Additionally, we further improve the results for bridging anaphora resolution reported in Hou (2018) by combining our simple deterministic approach with Hou et al. (2013b)’s best system *MLN II*.

1 Introduction

Anaphora plays a major role in discourse comprehension and accounts for the coherence of a text. In contrast to *identity anaphora* which indicates that a noun phrase refers back to the same entity introduced by previous descriptions in the discourse, *bridging anaphora* or *associative anaphora* links anaphors and antecedents via lexico-semantic, frame or encyclopedic relations. *Bridging resolution* is the task to recognize bridging anaphors (e.g., **distribution arrangements** in

Example 1¹) and find links to their antecedents (e.g., *dialysis products* in Example 1).

(1) While the discussions between Delmed and National Medical Care have been discontinued, Delmed will continue to supply *dialysis products* through National Medical after their exclusive agreement ends in March 1990, Delmed said. In addition, Delmed is exploring **distribution arrangements** with Fresenius USA, Delmed said.

Most previous empirical research on bridging (Poesio and Vieira, 1998; Poesio et al., 2004; Markert et al., 2003; Lassalle and Denis, 2011; Hou et al., 2013b) focus on *bridging anaphora resolution*, a subtask of bridging resolution that aims to choose the antecedents for bridging anaphors. 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 counts using certain syntactic patterns. However, such patterns only consider head noun knowledge and hence are not sufficient for bridging relations which require the semantics of modification. In Example 1, in order to find the antecedent (*dialysis products*) for the bridging anaphor “**distribution arrangements**”, we have to understand the semantics of the modification “**distribution**”.

Over the past few years, word embeddings gained a lot popularity in the NLP community. State-of-the-art word vectors such as word2vec skip-gram (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

¹All examples, if not specified otherwise, are from IS-Notes (Markert et al., 2012). Bridging anaphors are typed in boldface, antecedents in italics throughout this paper.

comprehension (Chen et al., 2016) and coreference resolution (Lee et al., 2017).

Recently, Hou (2018) found that these vanilla word embeddings capture both “genuine” similarity and relatedness, and hence they are not suitable for bridging anaphora resolution which requires lexical association knowledge instead of semantic similarity information between synonyms or hypernyms. Hou (2018) created word embeddings for bridging (*embeddings_PP*) by exploring the syntactic structure of noun phrases (NPs) to derive contexts for nouns in the GloVe model.

However, *embeddings_PP* only contains the word representations for nouns. In this paper, we improve *embeddings_PP* by combining it with *GloVe*. The resulting word embeddings (*embeddings_bridging*) are a more general lexical knowledge resource for bridging anaphora resolution. Compared to *embeddings_PP*, the coverage of lexicon in *embeddings_bridging* is much larger. Also the word representations for nouns without the suffix “_PP” are more accurate because they are trained on many more instances in the vanilla *GloVe*. Based on this general vector space, we develop a deterministic algorithm to select antecedents for bridging anaphors. Our approach combines the semantics of an NP’s head with the semantics of its modifications by vector average using *embeddings_bridging*. We show that this simple, efficient method achieves the competitive results on ISNotes for the task of bridging anaphora resolution compared to the best system in Hou et al. (2013b) which explores Markov Logic Networks to model the problem.

The main contributions of our work are: (1) a general word representation resource² for bridging; and (2) a simple yet competitive deterministic approach for bridging anaphora resolution which models the meaning of an NP based on its head noun and modifications.

2 Related Work

Lexical/world knowledge for bridging: Hou et al. (2013a) explored various lexico-semantic features for bridging anaphora recognition. Hou (2016) proposed an attention-based LSTM model with pre-trained word embeddings for information status classification and reported moderate results for bridging recognition. Previous work on bridg-

ing anaphora resolution (Poesio et al., 2004; Lassalle and Denis, 2011; Hou et al., 2013b) explored word co-occurrence counts in certain syntactic preposition patterns to calculate word relatedness. For instance, the big hit counts of the query “*the door of the house*” in large corpora could indicate that *door* and *house* stand in a part-of relation. These patterns encode associative relations between nouns which cover a variety of bridging relations. Unlike previous work which only consider a small number of prepositions per anaphor, the PP context model (Hou, 2018) uses all prepositions for all nouns in big corpora. It also includes the possessive structure of NPs. In this paper, we further improve Hou (2018)’s *embeddings_PP* by combining it with the vanilla *GloVe*. The resulting word embeddings (*embeddings_bridging*) are a more general lexical knowledge resource for bridging resolution. In addition, it enables efficient computation of word association strength through low-dimensional matrix operations.

Bridging anaphora resolution: regarding the methods to select antecedents for bridging anaphors, Poesio et al. (2004) applied a pairwise model combining lexical semantic features as well as salience features to perform mereological bridging resolution in the GNOME corpus³. To address the data sparseness problem (e.g., some part-of relations are not covered by WordNet), they used the Web to estimate the part-of relations expressed by certain syntactic constructions. Based on the method proposed by Poesio et al. (2004), Lassalle and Denis (2011) developed a system that resolves mereological bridging anaphors in French. The system was enriched with meronymic information extracted from raw texts. Such information was extracted in a bootstrapping fashion by iteratively collecting meronymic pairs and the corresponding syntactic patterns. Lassalle and Denis (2011) evaluated their system on mereological bridging anaphors annotated in the DEDE corpus and reported an accuracy of 23%.

Markert et al. (2012) released a corpus called ISNotes which contains unrestricted bridging annotations. Based on this corpus, Hou et al. (2013b) proposed a joint inference framework for bridging anaphora resolution using Markov logic networks (Domingos and Lowd, 2009). The framework resolves all bridging anaphors in one document together by modeling that semantically re-

²*embeddings_bridging* can be downloaded from <https://doi.org/10.5281/zenodo.1403164>

³The GNOME corpus is not publicly available.

Noun Phrases	Extracted Noun Pairs
<i>travelers in the train station</i>	<i>travelers_PP – station</i>
<i>travelers from the airport</i>	<i>travelers_PP – airport</i>
<i>hotels for travelers</i>	<i>hotels_PP – travelers</i>
<i>the destination for travelers</i>	<i>destination_PP – travelers</i>
<i>the company’s new appointed chairman</i>	<i>chairman_PP – company</i>

Table 1: Examples of noun phrases as well as the extracted noun pairs in *embeddings_PP*. Bold indicates the head noun of an NP.

lated anaphors are likely to share the same antecedent.

ISNotes is a challenging corpus for bridging. First, bridging anaphors are not limited to definite NPs as in previous work (Poesio et al., 1997, 2004; Lassalle and Denis, 2011). Also in ISNotes, the semantic relations between anaphor and antecedent are not restricted to meronymic relations. We therefore choose ISNotes to evaluate our algorithm for bridging anaphora resolution. Our approach is deterministic and simple, but achieves the competitive results compared to the advanced machine learning-based approach (Hou et al., 2013b). We also improve the result reported in Hou (2018) on the same corpus by combining our deterministic approach with the best system from Hou et al. (2013b).

Just recently, two new corpora (Rösiger, 2018a; Poesio et al., 2018) with bridging annotations have become available and we notice that the definitions of bridging in these corpora are different from the bridging definition in ISNotes. We apply our algorithm with small adaptations to select antecedents for bridging anaphors on these corpora. The moderate results demonstrate that *embeddings_bridging* is a general word representation resource for bridging.

3 Word Representations for Bridging

3.1 Word Embeddings Based on PP Contexts (*embeddings_PP*)

We briefly describe Hou (2018)’s *embeddings_PP* in this section. *embeddings_PP* released by Hou (2018) contains 100-dimensional vectors for 276k nouns. It is trained over 197 million noun pairs extracted from the automatically parsed Gigaword corpus (Parker et al., 2011; Napoles et al., 2012). The author generates these noun pairs by exploring the syntactic prepositional and possessive structures of noun phrases. These two structures encode a variety of bridging relations between

anaphors and their antecedents. For instance, the prepositional structure in “*the door of the house*” indicates the part-of relation between “*door*” and “*house*”. More specifically, for NPs containing the prepositional structure (e.g., **X** *preposition* **Y**) or the possessive structure (e.g., **Y** ’s **X**), the author extracts the noun pair “**X_PP–Y**”. Note that the head of the NP is always on the left and the noun modifier is always on the right. In addition, the suffix “_PP” is added for the nouns on the left. Table 1 shows a few examples of noun phrases together with the extracted noun pairs.

Hou (2018) showed that the suffix “_PP” plays an important role for the model to learn the asymmetric relations between the head nouns and their noun modifiers from the extracted noun pairs. For instance, among the top five nearest neighbors in *embeddings_PP*, “*president_PP*” is mostly related to countries or organizations (e.g., “*federation*”, “*republic*”, or “*USA*”), while “*president*” is mostly related to words which have the same semantic type as “*president*” (e.g., “*minister*”, “*mayor*”, or “*governor*”).

3.2 Word Representations for Bridging (*embeddings_bridging*)

embeddings_PP described in the previous section only contains word representations for nouns. To improve the coverage of lexical information, we create a general word representation resource *embeddings_bridging* by merging *embeddings_PP* with the original GloVe vectors trained on Gigaword and Wikipedia datasets. Specifically, given the 100 dimension word embeddings *embeddings_PP* and *GloVe*, we first create a 100 dimension vector *v_filler* with the value of each dimension as 0.1⁴. Let v_{1_w} represent the vector for the word w in *GloVe*, v_{2_w} represent the vector for the word w in *embeddings_PP*, if a word w appears

⁴Theoretically, any 100 dimension random vector with uniform distribution could be used as *v_filler*.

	<i>embeddings_bridging</i>	<i>embeddings_PP</i>
dimension size	200	100
vocabulary size	532,768	276,326
word type	all	nouns

Table 2: Comparison between *embeddings_bridging* and *embeddings_PP*

Category	Relation	prototypical word pair example	<i>embeddings_bridging</i>	<i>embeddings_PP</i>	GloVe
<i>PART-WHOLE</i>	<i>Object: Component</i>	{face: nose}	0.43	0.27	0.40
<i>PART-WHOLE</i>	<i>Event: Feature</i>	{wedding: bride}	0.46	0.29	0.15
<i>PART-WHOLE</i>	<i>Creature: Possession</i>	{author: copyright}	0.23	0.14	0.21
<i>PART-WHOLE</i>	<i>Activity: Stage</i>	{buying: shopping}	0.32	0.30	0.25

Table 3: Spearman’s rank correlation coefficient ρ for typical *PART-WHOLE* bridging relations using *embeddings_bridging*, *embeddings_PP* and *GloVe*.

both in *GloVe* and in *embeddings_PP*, its vector in *embeddings_bridging* is the concatenation of v_{1_w} and v_{2_w} . For the word w_1 which only appears in *GloVe*, its vector in *embeddings_bridging* is the concatenation of $v_{1_{w_1}}$ and v_{filler} . Finally, for the word w_2 which only appears in *embeddings_PP* (all the words with the suffix “_PP”), we construct its vector by concatenating v_{filler} and $v_{2_{w_2}}$. The resulting 200 dimension word embeddings (*embeddings_bridging*) is a general lexical resource for bridging. Table 2 compares the main features between *embeddings_bridging* and *embeddings_PP*. In the next section, we will compare *embeddings_bridging* with *embeddings_PP* and the original *GloVe* on a few typical bridging relations in the task of measuring relational similarity (Jurgens et al., 2012). Moreover, in Section 5.3 and Section 5.4, we show that using *embeddings_bridging* yields better results than using *embeddings_PP* for bridging anaphora resolution.

3.3 Measuring Relational Similarity on Typical Bridging Relations

We evaluate our *embeddings_bridging* quantitatively using a few typical bridging relations from SemEval-2012 Task 2 (Jurgens et al., 2012). The shared task aims to rank word pairs by the degree to which they are prototypical members of a given relation class. For instance, given the prototypical word pairs {*wedding–bride*, *rodeo–cowboy*, *banquet–food*} for the relation *Event:Feature*, we would like to know among the input word pairs {*school–students*, *circus–clown*, *meal–food*, *lion–zoo*}, which one represents the relation best.

SemEval-2012 Task 2 contains 79 relation classes chosen from Bejar et al. (1991). These relations fall into ten main categories, including *SIMILAR*, *PART-WHOLE*, *CONTRAST* and more. Each relation class is paired with a few prototypical word pairs and a list of around 40 word pairs which are ranked by humans according to their degree of similarity to the corresponding relation. We choose all typical bridging relations under the *PART-WHOLE* category and evaluate our *embeddings_bridging* in terms of ranking the list of word pairs for each relation. Spearman’s rank correlation coefficient ρ is used to evaluate a system by comparing the system’s ranking of the word pairs against the gold standard ranking.

Following Zhila et al. (2013), we calculate the relational similarity between word pairs using cosine similarity. Let (w_1, w_2) and (w_3, w_4) be the two word pairs, v_1, v_2, v_3, v_4 be the corresponding vectors for these words. We first normalize all word vectors to unit vectors, then the relational similarity between (w_1, w_2) and (w_3, w_4) is calculated as:

$$\frac{(v_1 - v_2) \cdot (v_3 - v_4)}{\|v_1 - v_2\| \|v_3 - v_4\|} \quad (1)$$

For each chosen relation class, we rank the list of word pairs according to their mean relational similarity to the given prototypical word pairs. Table 3 shows the results of Spearman’s rank correlation coefficient ρ for each typical bridging relation using *embeddings_bridging*, *embeddings_PP*, and *GloVe*, respectively. Note that when using *embeddings_bridging* and *embeddings_PP*, we add the suffix “_PP” to the potential bridging anaphor for

bridging anaphor: distribution arrangements				
Ante. Candidates	Head	Head + Modifiers	$dist_h$	$dist_{hm}$
<i>the discussions between Delmed and National Medical Care</i>	{ <i>discussions</i> }	{ <i>discussions</i> }	0.05	-0.10
<i>Delmed</i>	{ <i>delmed</i> }	{ <i>delmed</i> }	—	—
<i>National Medical Care</i>	{ <i>care</i> }	{ <i>care</i> }	0.08	0.10
<i>dialysis products</i>	{ <i>products</i> }	{ <i>dialysis, products</i> }	0.06	0.17
<i>National Medical</i>	{ <i>medical</i> }	{ <i>medical</i> }	0.02	-0.01
<i>their</i>	{ <i>their</i> }	{ <i>their</i> }	-0.05	-0.01
<i>their exclusive agreement</i>	{ <i>agreement</i> }	{ <i>exclusive, agreement</i> }	0.07	0.03

Table 4: The cosine similarities between the bridging anaphor **distribution arrangements** and its antecedent candidates for Example 1. $dist_h$ indicates the cosine similarity between {**arrangements_PP**} and the candidate head, $dist_{hm}$ the cosine similarity between {**distribution_PP, arrangements_PP**} and Head+Modifiers. “—” means *Delmed* is not present in *embeddings_bridging* and therefore we neglect this candidate.

each word pair (e.g., {wedding: bride_PP}). As shown in Table 3, using *embeddings_bridging* performs better than both using *embeddings_PP* and using the vanilla GloVe vectors on these four part-of relation classes. This partially indicates that *embeddings_bridging* could capture lexical knowledge for bridging relations.

4 A Deterministic Algorithm for Bridging Anaphora Resolution

In this section, we describe our deterministic algorithm based on *embeddings_bridging* for bridging anaphora resolution. For each anaphor a , we 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. We exclude an NP from E_a if it is a bridging anaphor because a bridging anaphor is rarely to be an antecedent for another bridging anaphor. We also exclude NPs whose semantic types are “time” from E_a if a is not a time expression. This is because time expressions are related to a lot of words in the corpus in which we learned *embeddings_bridging* from. Therefore we only keep them as the antecedent candidates for bridging anaphors whose semantic types are “time” (see Example 2).

(2) As a presidential candidate in 1980, George Bush forthrightly expressed his position on abor-

tion in an interview with Rolling Stone magazine published **that March**.

Given an anaphor a and its antecedent candidate list E_a , we predict the most semantically related NP among all NPs in E_a as the antecedent for a . In case of a tie, the closest one is chosen to be the predicted antecedent.

The relatedness is measured via cosine similarity between the vector representation of the anaphor and the vector representation of the candidate. More specifically, given a noun phrase np_1 , we first construct a list N which consists of the head and all common nouns (e.g., *earthquake victims*), adjectives (e.g., *economical sanctions*), and ed/ing participles (e.g., *the collapsed roadway* and *the landing site*) appearing before the head. If np_1 contains a post-modifier NP np_2 via the preposition “of”, we also add the above premodifiers and the head of np_2 to the list N (e.g., *the policies of racial segregation*). Finally, the noun phrase np_1 is represented as a vector v using the following formula, where the suffix “_PP” is added to each n if np_1 is a bridging anaphor and its semantic type is not time:

$$v = \frac{\sum_{n \in N} \text{embeddings_bridging}_n}{|N|} \quad (2)$$

The underlying intuition of adding NP modifications to the list N is that the above mentioned modifiers also represent core semantics of an NP, therefore we should consider them when selecting antecedents for bridging anaphors. For instance, as shown in Table 4, for Example 1, the cosine similarity between {**arrangements_PP**} and

{*products*} is 0.06, while the cosine similarity between {**distribution_PP**, **arrangements_PP**} and {*dialysis*, *products*} is 0.17.

If none of the words in N is present in *embeddings_bridging*, we simply neglect the noun phrase *np1*. Note that we do not add the suffix “_PP” to a bridging anaphor representing time information, because such an anaphor is likely to have the same semantic type antecedent (see Example 2). Therefore we use semantic similarity instead of relatedness to find its antecedent.

5 Experiments

5.1 Dataset

For the task of bridging anaphora resolution, we use the dataset ISNotes⁵ released by Markert et al. (2012). 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). As stated in Section 2, bridging anaphors in ISNotes are not limited to definite NPs as in previous work (Poesio et al., 1997, 2004; Lassalle and Denis, 2011). 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. This includes encyclopedic or frame relations such as *restaurant* – **the waiter** as well as context-specific relations such as *palms* – **the thieves**. In Example 1, “*dialysis products*” is the “theme” of the **distribution arrangements**. More specifically, “*dialysis products*” belongs to the frame element “Individuals” in the “Dispersal” frame that is triggered by “**distribution arrangements**”.

5.2 Experimental Setup

Following Hou et al. (2013b)’s experimental setup, we resolve bridging anaphors to entity antecedents. Entity information is based on the OntoNotes coreference annotation. We also use the OntoNotes named entity annotation to assign NPs the semantic type “*time*” if their entity types are “date” or “time”.

⁵<http://www.h-its.org/en/research/nlp/isnotes-corpus>

In Hou et al. (2013b), features are extracted by using entity information. For instance, the raw hit counts of the preposition pattern query (e.g., **arrangements** of *products*) for a bridging anaphor a and its antecedent candidate e is the maximum count among all instantiations of e . In our experiments, we simply extend the list of antecedent candidates E_a (described in Section 4) to include all instantiations of the original entities in E_a . Note that our simple antecedent candidate selection strategy (described in Section 4) allows us to include 76% of NP antecedents compared to 77% in *pairwise model III* from Hou et al. (2013b) where they add top 10% salient entities as additional antecedent candidates. In Hou et al. (2013b), salient entities on each text are measured through the lengths of the coreference chains based on the gold coreference annotation.

Following Hou et al. (2013b), we measure *accuracy* on the number of bridging anaphors, instead of on all links between bridging anaphors and their antecedent instantiations. We calculate how many bridging anaphors are correctly resolved among all bridging anaphors.

5.3 Using NP Head Alone

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 ⁶. The relatedness is measured via cosine similarity between the head of the anaphor (plus the postfix “_PP” if the anaphor is not a time expression) and the head of the candidate. We run experiments on the following four word embeddings: the original GloVe vectors trained on Gigaword and Wikipedia 2014 dump (*GloVe_GigaWiki14*), GloVe vectors that we trained on Gigaword only (*GloVe_Giga*), word vectors from Hou (2018) (*embeddings_PP*), and our word representation resource described in Section 3.2 (*embeddings_bridging*). Note that for the first two word vectors, we do not add the suffix “_PP” to the anaphor’s head since such words do not exist in *GloVe_GigaWiki14* and *GloVe_Giga*.

Table 5 lists the results for bridging anaphora resolution based on different word representation resources⁷. We notice that there is not

⁶In case of a tie, the closest one is chosen to be the predicted antecedent.

⁷Note that the results for the first three word embeddings are slight better than the ones reported in Hou (2018). This is due to the improved antecedent candidate selection strategy described in Section 4.

	acc
<i>GloVe_GigaWiki14</i>	21.42
<i>GloVe_Giga</i>	21.87
<i>embeddings_PP</i>	33.03
<i>embeddings_bridging</i>	34.84

Table 5: Results of using NP head alone for bridging anaphora resolution based on different word representation resources. Bold indicates statistically significant differences over the baselines (two-sided paired approximate randomization test, $p < 0.01$).

much difference between *GloVe_GigaWiki14* and *GloVe_Giga*. We find that using *embeddings_PP* achieves an accuracy of 33.03% on the ISNotes corpus, which outperforms the results based on *GloVe_GigaWiki14* and *GloVe_Giga* by a large margin. Using *embeddings_bridging* further improves the result by 1.8%. Although the improvement is not significant, we suspect that the representations for words without the suffix “_PP” in *embeddings_bridging* are more accurate because they are trained on many more instances in the vanilla GloVe vectors (*GloVe_GigaWiki14*).

5.4 Using NP Head + Modifiers

We carried out experiments using the deterministic algorithm described in Section 4 together with different word embeddings. Again we do not add the suffix “_PP” to the bridging anaphors for *GloVe_GigaWiki14* and *GloVe_Giga*.

Table 6 lists the best results of the two models for bridging anaphora resolution from Hou et al. (2013b). *pairwise model III* is a pairwise mention-entity model based on various semantic, syntactic and lexical features. *MLN model II* is a joint inference framework based on Markov logic networks (Domingos and Lowd, 2009). It models that semantically or syntactically related anaphors are likely to share the same antecedent and achieves an accuracy of 41.32% on the ISNotes corpus.

The results for *GloVe_GigaWiki14* and *GloVe_Giga* are similar on two settings (*using NP head* vs. *using NP head + modifiers*). For *embeddings_PP*, the result on *using NP head + modifiers* (31.67%) is worse than the result on *using NP head* (33.03%). However, if we apply *embeddings_PP* to a bridging anaphor’s head and modifiers, and only apply *embeddings_PP* to the head noun of an antecedent candidate, we get an accuracy of 34.53%. Although the differences are not significant, it confirms that the information

	acc
models from Hou et al. (2013b)	
<i>pairwise model III</i>	36.35
<i>MLN model II</i>	41.32
NP head + modifiers	
<i>GloVe_GigaWiki14</i>	20.52
<i>GloVe_Giga</i>	20.81
<i>embeddings_PP</i>	31.67
<i>embeddings_bridging</i>	39.52

Table 6: Results of using NP head plus modifications in different word representations for bridging anaphora resolution compared to the best results of two models from Hou et al. (2013b). Bold indicates statistically significant differences over the other models (two-sided paired approximate randomization test, $p < 0.01$).

from the modifiers of the antecedent candidates in *embeddings_PP* hurts the performance. This corresponds to our observations in the previous section that the representations for words without the suffix “_PP” in *embeddings_PP* are not as good as in *embeddings_bridging* due to less training instances.

Finally, our method based on *embeddings_bridging* achieves an accuracy of 39.52%, which is competitive to the best result (41.32%) reported in Hou et al. (2013b). There is no significant difference between *NP head + modifiers* based on *embeddings_bridging* and *MLN model II* (randomization test with $p < 0.01$).

To gain an insight into the contribution of *embeddings_bridging* on different relation types, we analyze the results of our method using *embeddings_bridging* on three relation types: *set-of*, *part-of*, and *other*. The accuracies on these three relation types are 17.78%, 50.0%, and 39.16%, respectively. This suggests that in the future we should include more context for bridging anaphors that hold the *set-of* relation to their antecedents, because the head nouns of such anaphors often do not bear any specific meanings (e.g., **Another**).

5.5 Analysis of Modifiers

To better understand the role of NP modifiers in our method, we carried out experiments on *embeddings_bridging* using different set of modifiers (see Table 7). It seems that among all three types of modifiers, compared to using NP head alone, adding noun modifiers has the positive impact (36.65% on *NP head + noun modifiers* vs. 34.84% on *NP head*). Although adding only ad-

<i>embeddings_bridging</i>	acc
<i>NP head</i>	34.84
+ <i>all modifiers</i>	39.52
+ <i>noun modifiers</i>	36.65
+ <i>adjective modifiers</i>	34.84
+ <i>ed/ing participle modifiers</i>	34.84
+ <i>noun&adjective modifiers</i>	38.31
+ <i>noun&ed/ing participle modifiers</i>	36.80
+ <i>adjective&ed/ing participle modifiers</i>	34.84

Table 7: Results of using NP head plus different modifications in *embeddings_bridging*.

jective modifiers does not have influence on results, combining them with noun modifiers yields some improvement over adding only noun modifiers (38.31% on *NP head + noun&adjective modifiers* vs. 36.65% on *NP head + noun modifiers*). On the other hand, *ed/ing* participle modifiers only have a small positive impact over *NP head + noun modifiers* when combining with noun modifiers.

5.6 Combining NP Head + Modifiers with MLN II

For bridging anaphora resolution, Hou (2018) integrates a much simpler deterministic approach by combining an NP head with its noun modifiers (appearing before the head) based on *embeddings_PP* into the *MLN II* system (Hou et al., 2013b). Similarly, we add a constraint on top of *MLN II* using our deterministic approach (*NP head + modifiers*) based on *embeddings_bridging*. Table 8 lists the results of different systems⁸ for bridging anaphora resolution in ISNotes. It shows that combining our deterministic approach (*NP Head + modifiers*) with *MLN II* slightly improves the result compared to Hou (2018).

Although combining *NP Head + modifiers* with *MLN II* achieves significant improvement over *NP Head + modifiers*, we think the latter has its own value. Our deterministic algorithm is simpler and more efficient compared to *MLN model II + em-*

⁸We also reimplement the algorithms from Schulte im Walde (1998) and Poesio et al. (2004) as baselines (Table 8). Schulte im Walde (1998) resolved bridging anaphors to the closest antecedent candidate in a high-dimensional space. We use the 2,000 most frequent words (adjectives, common nouns, proper nouns, and lexical verbs) from Gigaword as the context words. Poesio et al. (2004) applied a pairwise model combining lexical semantic features and salience features to perform mereological bridging resolution in the GNOME corpus. We use a Naive Bayes classifier with standard settings in WEKA (Witten and Frank, 2005) and apply the best first strategy to select the antecedent for each anaphor.

beddings_bridging, which contains many complicated features and might be hard to migrate to other bridging corpora. Moreover, our algorithm is “unsupervised” and requires no training when applied to other English bridging corpora.

5.7 Resolving Bridging Anaphors in Other Corpora

Recently, two new corpora containing bridging annotation have become available. The BASHI corpus (Rösiger, 2018a) contains 459 bridging NPs and their antecedents in 50 World Street Journal articles. Similar to ISNotes, BASHI includes both definite and indefinite referential bridging anaphors. In addition, comparative anaphora is also considered as bridging anaphora in BASHI.

Another new corpus for bridging is the second release of the ARRAU corpus, which contains 5,512 bridging pairs in three different domains (Poesio et al., 2018). However, most bridging links in ARRAU are purely lexical bridging pairs, and only a small subset of the annotated pairs contains truly anaphoric bridging anaphors (Rösiger et al., 2018). Following Rösiger et al. (2018), we focus on resolving bridging anaphors in the news text domain (RST).

Based on *embeddings_bridging*, we apply our deterministic algorithm with small adaptations to resolve bridging anaphors to entity antecedents on the BASHI and ARRAU (RST) corpora. Specifically, for the BASHI corpus, we do not add NPs from the first sentence to the list of antecedent candidates E_a . This is because the phenomenon of globally salient antecedents being linked to all anaphors in a text is less obvious in BASHI. In addition, comparative anaphors often have the same semantic class as their antecedents, therefore we do not add the suffix “_PP” to a bridging anaphor if it is a comparative anaphor.

For the ARRAU corpus, we construct the list of antecedent candidates E_a using NPs preceding a from the same sentence as well as from the previous ten sentences. Since most bridging pairs in ARRAU are lexical bridging (e.g., *Tokyo – Japan, other nations – Britain*) and anaphors often have the same semantic type as their antecedents, we do not add the suffix “_PP” to bridging anaphors.

Table 9 lists the results of bridging anaphora resolution in the BASHI and ARRAU corpora, respectively. On the test set of the ARRAU (RST) corpus, Rösiger (2018b) proposed a modi-

	System	acc
Baselines	<i>Schulte im Walde (1998)</i>	13.68
	<i>Poesio et al. (2004)</i>	18.85
Models from Hou et al. (2013b)	<i>pairwise model III</i>	36.35
	<i>MLN model II</i>	41.32
Hou (2018)	<i>MLN model II + embeddings_PP (NP head + noun pre-modifiers)</i>	45.85
This work	<i>embeddings_bridging (NP head + modifiers)</i>	39.52
	<i>MLN model II + embeddings_bridging (NP head + modifiers)</i>	46.46

Table 8: Results of different systems for bridging anaphora resolution in ISNotes. Bold indicates statistically significant differences over the other models (two-sided paired approximate randomization test, $p < 0.01$).

Corpus	Bridging Type	# of Anaphors	acc
<i>BASHI</i>	referential, including comparative anaphora	452	27.43
<i>BASHI</i>	referential, excluding comparative anaphora	344	29.94
<i>ARRAU (RST Train)</i>	mostly lexical, some referential	2,325	31.44
<i>ARRAU (RST Test)</i>	mostly lexical, some referential	639	32.39

Table 9: Results of resolving bridging anaphors in other corpora. Number of bridging anaphors is reported after filtering out a few problematic cases on each corpus.

fied rule-based system based on Hou et al. (2014)’s work and reported an accuracy of 39.8% for bridging anaphora resolution. And our algorithm achieves an accuracy of 32.39% using only *embeddings_bridging*. Overall, the reasonable performance on these two corpora demonstrates that *embeddings_bridging* is a general word representation resource for bridging.

6 Conclusions

We improve the word representation resource *embeddings_PP* (Hou, 2018) by combining it with *GloVe*. The resulting word embeddings (*embeddings_bridging*) are a more general word representation resource for bridging. Based on *embeddings_bridging*, we propose a deterministic approach for choosing antecedents for bridging anaphors. We show that this simple and efficient method achieves the competitive result on bridging anaphora resolution compared to the advanced machine learning-based approach in Hou et al. (2013b) which is heavily dependent on a lot of carefully designed complex features. We also demonstrate that using *embeddings_bridging* yields better results than using *embeddings_PP* for bridging anaphora resolution.

For the task of bridging anaphora resolution, Hou et al. (2013b) pointed out that considering only head noun knowledge is not enough and future work needs to explore wider context to re-

solve context-specific bridging relations. In this work we explore the context within NPs—that is, we combine the semantics of certain modifications and the head by vector average using *embeddings_bridging*. But in some cases, knowledge about NPs themselves is not enough for resolving bridging. For instance, in Example 3, knowing that **any loosening** has the ability to “rekindle inflation” from the context of the second sentence can help us to find its antecedent “*the high rates*” (which is used to against inflation).

(3) Chancellor of the Exchequer Nigel Lawson views *the high rates* as his chief weapon against inflation, which was ignited by tax cuts and loose credit policies in 1986 and 1987. Officials fear that **any loosening** this year could rekindle inflation or further weaken the pound against other major currencies.

In the future, we will study how to integrate context outside of NPs for the task of choosing antecedents for bridging anaphors. Also we hope that our word representation resource will facilitate other related research problems such as semantic role labeling.

Acknowledgments

The author appreciates the valuable feedback from the anonymous reviewers and would like to thank Massimo Poesio for sharing the ARRAU corpus.

References

- I.I. Bejar, R. Chaffin, and S.E. Embretson. 1991. *Cognitive and psychometric analysis of analogical problem solving*. Springer-Verlag.
- 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.
- Pedro Domingos and Daniel Lowd. 2009. *Markov Logic: An Interface Layer for Artificial Intelligence*. Morgan Claypool Publishers.
- 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. 2018. Enhanced word representations for bridging anaphora resolution. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, New Orleans, Louisiana, 1–6 June 2018, pages 1–7.
- 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.
- 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.
- 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.
- David A. Jurgens, Peter D. Turney, Saif M. Mohammad, and Keith J. Holyoak. 2012. Semeval-2012 task 2: Measuring degrees of relational similarity. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics*, Montréal, Canada, 7–8 June 1999, pages 356–364.
- 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.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Copenhagen, Denmark, 7–11 November 2017, pages 188–197.
- 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.
- Katja Markert, Malvina Nissim, and Natalia N. Modjeska. 2003. Using the web for nominal anaphora resolution. In *Proceedings of the EACL Workshop on the Computational Treatment of Anaphora*. Budapest, Hungary, 14 April 2003, pages 39–46.
- 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, Yulia Grishina, Varada Kolhatkar, Nafise Sadat Moosavi, Ina Rösiger, Adam Roussel, Fabian Simonjetz, Alexandra Uma, Olga Uryupina, Juntao Yu, and Heike Zinsmeister. 2018. Anaphora resolution with the ARRAU corpus. In *Proceedings of the Workshop on Computational Models of Reference, Anaphora and Coreference*. New Orleans, Louisiana, June 6, 2018, pages 11–22.
- 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 and Renata Vieira. 1998. A corpus-based investigation of definite description use. *Computational Linguistics*, 24(2):183–216.
- 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.
- 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.
- Ina Rösiger. 2018a. BASHI: A corpus of wall street journal articles annotated with bridging links. In *Proceedings of the 11th International Conference on Language Resources and Evaluation*, Miyazaki, Japan, 7–12 May 2018, pages 382–388.
- Ina Rösiger. 2018b. Rule- and learning-based methods for bridging resolution in the ARRAU corpus. In *Proceedings of the Workshop on Computational Models of Reference, Anaphora and Coreference*. New Orleans, Louisiana, June 6, 2018, pages 23–33.
- Ina Rösiger, Arndt Riester, and Jonas Kuhn. 2018. Bridging resolution: Task definition, corpus resources and rule-based experiments. In *Proceedings of the 27th International Conference on Computational Linguistics*, Santa Fe, New-Mexico, USA, 20–26 August 2018, pages 3516–3528.
- Sabine Schulte im Walde. 1998. Resolving bridging descriptions in high-dimensional space. Master’s thesis, University of Edinburgh, Centre for Cognitive Science.
- 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.
- Ian H. Witten and Eibe Frank. 2005. *Data Mining: Practical Machine Learning Tools and Techniques*, 2nd edition. Morgan Kaufmann, San Francisco, Cal.
- Alisa Zhila, Scott Wen-tau Yih, Geoffrey Zweig, Chris Meek, and Tomas Mikolov. 2013. Combining heterogeneous models for measuring relational similarity. 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 1000–1009.