

Data Augmentation Methods for Anaphoric Zero Pronouns

Abdulrahman Aloraini^{1,2}, Massimo Poesio¹

¹Queen Mary University of London, United Kingdom

²Qassim University, Saudi Arabia

{a.aloraini, m.poesio}@qmul.ac.uk

Abstract

In pro-drop language like Arabic, Chinese, Italian, Japanese, Spanish, and many others, un-realized (null) arguments in certain syntactic positions can refer to a previously introduced entity, and are thus called anaphoric zero pronouns. The existing resources for studying anaphoric zero pronoun interpretation are however still limited. In this paper, we use five data augmentation methods to generate and detect anaphoric zero pronouns automatically. We use the augmented data as additional training materials for two anaphoric zero pronoun systems for Arabic. Our experimental results show that data augmentation improves the performance of the two systems, surpassing the state-of-the-art results.

1 Introduction

In pronoun-dropping (pro-drop) languages such as Arabic (Eid, 1983), Chinese (Li and Thompson, 1979), Italian (Di Eugenio, 1990) and other romance languages (e.g., Portuguese, Spanish), Japanese (Kameyama, 1985), and others (Kim, 2000), certain arguments can be omitted in which a pronoun is used in English: such arguments are called zero-pronouns (ZP) in NLP. Anaphoric zero pronouns (AZP) are a special case of zero pronouns. AZPs refer to one or more noun entities previously mentioned in a text, as in the following example from the Arabic portion of OntoNotes (in Arabic OntoNotes 5.0, ZPs are denoted as * in Arabic text).

.. المفارقة الأخرى عن بوش هي عدم حماسه
للمؤتمر الدولي، إذ أنه من البداية، يريد * اجتماعا
مختلفا....

Ironically, Bush did not show any enthusiasm for the international conference, because since the beginning, (he) wanted to attend another conference ...

In the example, the AZP, indicated with '*', refers to a singular masculine person that has been mentioned previously, Bush (بوش). English is not a pro-drop language (White, 1985) so the null argument in Arabic has to be expressed using an overt pronoun (he).

AZP resolution involves two tasks: identification and resolution. AZP identification is the task of identifying the locations where the pronoun was omitted (the * in the example). AZP resolution is the task of finding the antecedent that an AZP refers to (Bush in the example). Contextual information is required for these two tasks, such as that provided by the verb which has the AZP as an argument (in the example the verb is "wanted/يريد").

The AZP problem has inspired much research because AZP interpretation benefits many natural language processing (NLP) tasks such as machine translation (Mitkov and Schmidt, 1998). Although AZPs are common (Chen and Ng, 2016), they are not always annotated in NLP corpora. There are two reasons for this. First, AZPs have no surface realization and the focus in coreference is usually on arguments realized on the surface (Lee et al., 2005). Second, pro-drop languages have several types of ZPs, not all of which anaphoric, which can make it challenging to identify AZPs. Therefore, the number of datasets with annotated AZPs is small (Konno et al., 2020). We investigate therefore automatic methods for augmenting these existing data, applying them to Arabic. Our contributions are as follows:

- We apply various data augmentation methods to detect potential AZPs in unannotated sentences, and to generate sentences containing AZPs.
- We propose a method to automatically find the true antecedent of AZPs.
- The augmented data improve AZP identification and resolution for Arabic, and surpass the

current state-of-the-art results.

The rest of the paper is organized as follows. We discuss AZP and data augmentation related literature in Section 2. We explain our data-augmentation methods in Section 3. We discuss the evaluation settings and results in Section 4 and Section 5 respectively. We conclude in Section 6.

2 Related Work

2.1 Anaphoric Zero Pronouns

Arabic: There have been a few proposals devoted to AZP tasks and null arguments in general. [Green et al. \(2009\)](#) proposed a conditional-random-field (CRF) sequence classifier to detect Arabic noun phrases, and captured ZPs implicitly. [Bakr et al. \(2009\)](#) applied a statistical approach to detect empty categories. [Gabbard \(2010\)](#) proposed a pipeline made of maximum entropy classifiers which jointly make a CRF to retrieve Arabic empty categories. [Aloraini and Poesio \(2020b\)](#) proposed the first neural model for resolving Arabic AZP, but they did not consider the AZP identification step. [Aloraini and Poesio \(2020a\)](#) showed a multilingual approach to detect AZPs for Arabic and Chinese using BERT ([Devlin et al., 2018](#)).

Other languages: for Chinese, [Converse \(2006\)](#) studied AZP resolution and applied a rule-based approach that employed Hobbs algorithm ([Hobbs, 1978](#)) to resolve ZPs in the Chinese Treebank; however, did not attempt to automatically identify AZP. [Yeh and Chen \(2006\)](#) is another rule-based approach for Chinese which used a set of hand-engineered rules to identify and resolve AZPs. [Zhao and Ng \(2007\)](#) proposed the first machine learning approach to Chinese AZPs identification and resolution. They applied decision trees incorporated with a set of syntactic and positional features. Other proposals followed the machine learning approach targeting Chinese ([Kong and Zhou, 2010](#); [Chen and Ng, 2013, 2014, 2015, 2016](#); [Yin et al., 2016, 2017](#); [Liu et al., 2017](#); [Chang et al., 2017](#); [Yin et al., 2018](#); [Kong et al., 2019](#)). There has been also a great deal of research on identification and resolution of AZPs, particularly in Japanese ([Yoshimoto, 1988](#); [Kim and Ehara., 1995](#); [Aone and Bennett, 1995](#); [Seki et al., 2002](#); [Isozaki and Hirao, 2003](#); [Iida et al., 2006, 2007](#); [Sasano et al., 2008, 2009](#); [Sasano and Kurohashi, 2011](#); [Yoshikawa et al., 2011](#); [Hangyo et al., 2013](#); [Iida et al., 2015](#); [Yoshino et al., 2013](#); [Yamashiro et al.,](#)

[2018](#)), but also in other languages, including Korean ([Han, 2004](#); [Byron et al., 2006](#)), Spanish ([Fernández and Peral, 2000](#); [Rello and Ilisei, 2009](#)), Portuguese ([Rello et al., 2012](#)), Romanian ([Mihăilă et al., 2011](#)), Bulgarian ([Grigorova, 2013](#)), and Sanskrit ([Gopal and Jha, 2017](#)). [Iida and Poesio \(2011\)](#) proposed the first multilingual approach for AZP resolution.

These proposals applied their AZP systems to labeled dataset. Even though data augmentation improved many NLP tasks, as discussed in Section 2.2, none of the above proposals considered augmenting AZPs datasets automatically. As far as we know, ([Konno et al., 2020](#)) is the only proposal to use data augmentation for AZP resolution. Konno et al applied a data augmentation method called contextual data augmentation (CDA) to Japanese AZPs. CDA method is based on language models to generate different variants of a labeled input by masking certain tokens and then predict them. Our work is different in that we apply various methods to detect, generate, and resolve AZPs while they focus generating samples for AZP resolution. The generated AZP samples are of numerous types and they are automatically prepared for both AZP tasks.

2.2 Data Augmentation

Data augmentation is an active research topic and has been applied in different areas of research of NLP ([Feng et al., 2021](#); [Chen et al., 2021](#)). However, there has been very limited proposals for AZP data augmentation. [Zhang et al. \(2015\)](#) augmented data by replacing a word with its synonyms to improve text classification. [Sennrich et al. \(2015\)](#) augmented data by translating a sequence from one language to another, and then translating the sequence back into the original language. The new data were used to enhance the performance of neural machine translation models. [Wang et al. \(2018\)](#) examined various methods for data augmentation and proposed to randomly replace words in source and target languages to improve neural machine translation. [Şahin and Steedman \(2019\)](#) removed dependency links and modified tree nodes to create an augmented dataset for Part-of-Speech tagging. [Gulordava et al. \(2018\)](#) replaced words with other words that share the same Part-of-Speech, morphological features, and dependency labels, to improve subject-verb agreement models. [Feng et al. \(2019\)](#) introduced a pipeline called SMERTI which combines various data augmentation methods, such as, entity replacement,

similarity masking, and text in-filling. Grundkiewicz et al. (2019) used a spellchecker to augment training data which are then used to pre-train sequence-to-sequence models. Singh et al. (2019) introduced a cross-lingual data augmentation called XLDA, evaluated on 14 languages on a natural language inference (XNLI) benchmark and question-answering task. XLDA replaces segments of an inputs text with its translation in other languages. Kumar et al. (2019) proposed DiPS, a model that generates various paraphrased sentences used to train conversational agents and in text summarization tasks. Andreas (2019) proposed rule-based data augmentation, which replaces segments of inputs that share similar context to improve the training of n-gram and sequence-to-sequence language models. Guo et al. (2020) proposed a statistical approach called SeqMix to decide which token to use at each position of an input, and they also provided a framework that combines several data augmentation approaches for several NLP tasks. Chen et al. (2020) represented datasets as graphs and proposed methods to augment data based on graph theory for paraphrasing. Ding et al. (2020) trained a language model on the linearized version of the input to synthesize data for low-resource sequence-tagging. Feng et al. (2020) inserted character-level synthetic noise and word hypernyms to augment data for text generation. Louvan and Magnini (2020) used a set of augmentation methods that span words and modify sentences for slot filling and intent detection. (Ri et al., 2021) proposed a method to augment ZPs (not AZPs). They delete personal overt pronouns which results in extra sentences with no pronouns, and use them in Japanese-English translation. Other proposals investigated data augmentation for neural machine translation (Vaibhav et al., 2019; Gao et al., 2019; Nguyen et al., 2019), text classification (Kobayashi, 2018; Wei and Zou, 2019; Anaby-Tavor et al., 2020; Jindal et al., 2020), and question-answering (Kafle et al., 2017; Yang et al., 2019).

Data augmentation methods helped to improve various NLP tasks; however, very few proposals (Konno et al., 2020) considered the methods for AZP tasks which was focused on generating AZPs for resolution. This paper explores the direction of exploiting several data augmentation methods for AZP resolution and identification, and how they can benefit both tasks.

3 Methodology

We applied the following five methods to generate AZPs:

1. **OntoNotes Patterns (ONP):** AZPs may occur more frequently in the company of certain verbs. To find the most frequent collocations, we apply the t-test on the Part-of-Speech sequences of AZP sentences. We tried a window of one, two, and three and we found empirically a window of two to detect many correct AZP collocations.
2. **Removing Subject Mention (RSM):** AZPs are dropped subjects of verbs. When an (explicit) mention is the subject of a verb phrase, we remove the mention to obtain an AZP sentence.
3. **Masking Candidate Mentions (MCM),** also known as contextual data augmentation (CDA). We mask the true antecedents of AZPs and use a language model to find the semantically most similar words, which are then used to produce new sentences by replacing the original word.
4. **Back Translation (BT):** the AZP examples in OntoNotes are in Arabic. We translate them into English using the GoogleTrans API,¹ and then translate them back to Arabic.
5. **Changing Subject-verb Agreement (CSA):** the verb in an AZP construction and its true antecedent agree on number which can be in singular, dual, or plural form. We change their agreement number from one form to another. For example, if the AZP verb and its reference are in singular form, we change the verb and its reference to dual or plural form.

Table 1 shows the number of collected data of each method. AZPs interpretation involves two steps: AZP detection, and AZP resolution. The detection step finds if a sentence has AZPs while the resolution resolves AZPs to their true antecedent. The true antecedent might be present in the same AZP sentence or any previous sentences. In our experiments, we decided to collect AZP samples such that each sample has two sentences: the first has the true antecedent and the second has its AZP. While it is possible to automatically create AZP

¹<https://pypi.org/project/googletrans/>

Method	Training
OntoNotes Patterns (ONP)	369
Removing Subject (RSM)	1,196
Masking Candidate (MCM)	736
Back Translation (BT)	501
Changing Subject-Verb Agr. (CSA)	104
Total	2,906

Table 1: The number of the augmented data of each method, and their total.

samples of one sentence only (the AZP position and its true antecedent in the same sentence), but two-sentence samples provide more context and information about AZPs.

3.1 OntoNotes Patterns

AZP sentences may occur more frequently in certain constructions. To find these constructions, we apply a window of one, two, and three tokens on the Part-of-Speeches of AZP samples. We found a window of two (two words before an AZP and two words after as shown in Figure 1) to provide many correct collocations. We count the frequencies of the constructions and apply the t-test score as in (Manning and Schutze, 1999) to see how probable a construction is:

$$t = \frac{\bar{x} - \mu}{\sqrt{\frac{s^2}{N}}}$$

Where \bar{x} is the sample mean of a Part-of-Speech, and s^2 is its sample variance, N is its sample size, and μ is its distribution mean. We choose the highest five t-test scores and apply them to large public text datasets i.e Wikipedia. We tag the sentences in Wikipedia summary sections with their Part-of-Speech using CAMEL tools (Obeid et al., 2020) and apply the OntoNotes constructions on each sentence except the first. When one sentence matches a pattern, it suggests that the sentence might be an AZP. The AZP true antecedent is usually in the first sentence of the summary section. We join the first sentence and the detected sentence together to represent one AZP sample, as shown in Figure 3. In Section 3.6, we discuss in details why we follow this approach and the challenges of finding the true antecedent of an AZP.

3.2 Removing Subject Mentions

Subject nouns or pronouns of a verb phrase might be optional. Removing these subjects transforms a




Figure 1: An example of one construction of an AZP from OntoNotes. For every AZP position, we extract the POS of two words before and two after, in order to create one construction.

sentence into a null-subject, as shown in the following example:

تقع لندن على نهر التامز
London is located on the Thames river.

In the example, removing 'London' from the Arabic sentence changes it into an AZP sentence and grammatically correct. However, when the subject is removed, readers might not know what the AZP refers unless they have access to more context i.e previous sentences. Therefore, in our experiments each AZP sample has two sentences. The first sentence contains the mention that the AZP refers to, and the second contains the AZP gap.

3.3 Masking Candidate Mentions

We use the multilingual BERT² to mask antecedents of AZPs and replace them with their most similar semantic token. We replace the antecedent token with [MASK]³, and predict its most similar semantic word. We replace the original mention with the predicted mention from BERT. However, even though the predicted token might be semantically similar to the original token, it might not agree in number or gender with the AZP verb. For example, the antecedent in the original sample can be *student* and BERT replaces it with *students*. Such samples can distort the training for machine learning algorithms because it associates the morphemes with wrong morphological features. We address this problem in Section 3.7.

3.4 Back Translation

It has been shown that translating a sentence from one language to another, and then translating it back to the original language to be beneficial to some NLP tasks (Xie et al., 2018; Vaibhav et al., 2019;

²<https://github.com/google-research/bert>

³[MASK] is a special token in BERT which is used for prediction.

Zhang et al., 2019). Back Translation generates a paraphrased version of the original input adding noise, such as, semantic and syntactical changes. Therefore, we translate the Arabic samples to English, and translate them back to Arabic.

3.5 Changing Subject–verb Agreement

AZP verb and its antecedent usually agree on number, gender, and person. We change the verb and its antecedent number agreement from one form to another. For example, if the number of an AZP sample is singular, we change it to dual or plural. To find and generate verb inflections and mention number, we use CAMEL which consists of a set of NLP tools for Arabic.

The five methods we discussed are used create AZP samples, but they are applied differently. The first two (ONP and RSM) are used to detect AZPs on text while the other three (CSA, MCM, and BT) are used to generate AZP from existing AZP samples. Thus, we apply the two methods on the Wikipedia summary sections to initially collect AZPs, and then use the other three to generate extra samples. An illustration of all data augmentation methods and steps in Figure 2.

3.6 Finding the True Antecedent

An AZP refers to a preceding mention in a text. The mention can be in the same sentence as the AZP or in previous sentences. Finding the true mention in a sentence using patterns is not trivial because resolving AZPs involves reasoning, context, background knowledge of real world, and deep understanding of a language characteristic, such as, its morphology (Alnajadat, 2017). At the beginning of the experiment, we apply the two methods (ONP and RSM) on Arabic newspaper articles of Abu El-Khair corpus (El-Khair, 2017) to detect AZP locations. For resolving AZPs to their antecedents, we used AZP context features to find the true antecedent. The main features are the agreement morphemes between the AZP verb and its antecedent. The result was we managed to identify many AZP locations successfully; however, we encountered two problems when we tried finding their true antecedents. First, there were many mentions that share similar features as the true antecedent which led to many incorrect antecedent reference. Second, some locations of the true antecedents were very far from the AZP gap.

To alleviate the above-mentioned limitations of AZP resolution, we apply the two methods on

Wikipedia articles. Every Wikipedia articles focuses on a single topic especially in the summary section. Summary sentences can contain AZPs and they usually refer to the Wikipedia title. Therefore, we apply the five methods on the summary sentences, and we consider the title to be the true antecedent of every AZP. The true antecedent always appears in the first sentence i.e the title of Wikipedia page or part of it. So for every AZP in the other sentences of the same page, they refer to the mention in the first sentence. So each AZP sample consists of the first sentence of the Wikipedia page and the sentence where it has the AZP. The reason why we apply the methods only on Wikipedia summary sections because their sentences usually focus on briefly describing the entity of the page. The sentences of other sections might have AZPs but they also include other mentions which can make it challenging to resolve the AZPs to their true antecedent. An illustration of creating one sample from Wikipedia is in Figure 3. After collecting the samples using ONP and RSM methods, we apply the other methods MCM, BT, CSA to generate extra samples.

3.7 Filtering Samples

After data creation and generation, we found some samples to be grammatically incorrect or not AZPs. For example, sometimes MCM method replaces the original antecedent of an AZP with a mention of a different number or gender. BT method can be problematic in some cases because it predicts the dropped subject which converts an AZP sample into a non-AZP sample. Another problem of BT method is that the AZP verb in the original sample is used to track the AZP location. If the AZP verb is translated back into another similar verb, the AZP location can not be retrieved back. To enhance the quality of the augmented data produced by the MCM method, we use Camel tools⁴ which has a morphological analyser to find the gender and number of mentions. We use the analyser to see if the AZP verb and its antecedent agree on number and gender. For the grammatically incorrect cases, we remove samples that do not have complete agreements. For the added subject after the BackTranslation step, we remove the subjects in a similar way as we do in RSM method. If the verb of the original sample is absence after the BackTranslation, we remove the sample.

⁴https://github.com/CAMEL-Lab/camel_tools

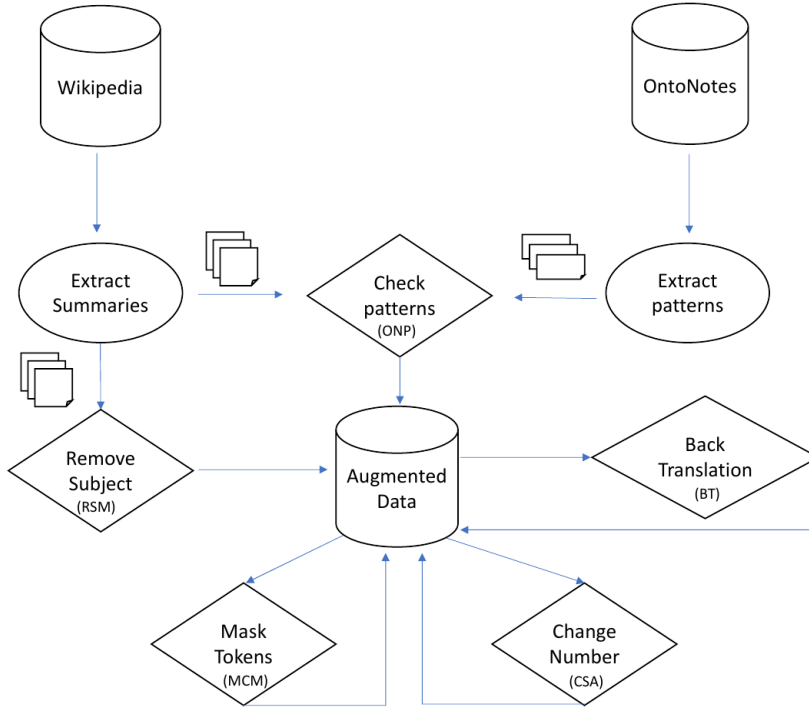


Figure 2: Overall picture of all data augmentation methods for Anaphoric Zero Pronouns. We extract patterns from OntoNotes and check them on Wikipedia summary sentences (ONP method). Removing Subject (RSM) is also applied to the summary sentences. after collecting AZP samples, We create extra data applying MCM, BT, and CSA methods.

4 Evaluation

4.1 Dataset

Most existing Arabic corpora do not have AZP annotated. As far as we know, OntoNotes 5.0 is the only source that has Arabic AZPs annotated. The distribution of documents, sentences, words, and AZPs in OntoNotes are in Table 2.

4.2 AZP systems

For AZP identification, we use the model by (Aloraini and Poesio, 2020a) which is a binary classifier that takes a candidate ZP location as input, and classifies whether it as an AZP or not. For AZP resolution, we use the model by (Aloraini and Poesio, 2020b) which combines BERT representations and additional task-related features of AZPS to learn their true antecedent. We adopt their systems and apply the same settings on Arabic AZPs. We then evaluate each method of our data augmentation for both systems to see its effect. We then try using all the collected data, and try various binary combinations as we will see in Section 5.

Category	Training	Dev	Test
Documents	359	44	44
Sentences	7,422	950	1,003
Words	264,589	30,942	30,935
AZPs	3,495	474	412

Table 2: Basic statistics about AZPs in Arabic OntoNotes. The total number of AZPs is 3495, 474, and 412 for training, development, and test respectively.

4.3 Data Preprocessing

Arabic is a morphologically rich language with a large set of morphological properties. Arabic text can suffer from sparsity (different forms for the same word) and ambiguity (same form for numerous words) if the text is not pre-processed. There are two reasons for these problems. First, certain letters can have different forms which are usually misspelled, such as the various forms of the letter “alif”. Second, same word can have a fully diacritized, partially diacritized or undiacritized forms. Retaining diacritics can complex word representation and model training (Habash and Sadat, 2006). Pre-processing Arabic text improves the overall performance. For example, Aloraini et al. (2020) fol-

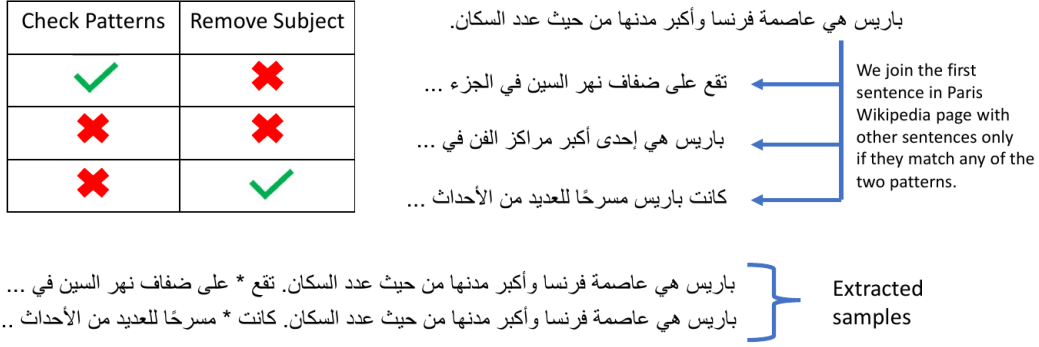


Figure 3: Illustration of finding AZPs using OntoNotes Patterns (ONP) and Removing Subject Mention (RSM) on Paris, from Arabic Wikipedia. Given a Wikipedia page, we join the first sentence with any other sentence in the summary only if they match any of the two data augmentation methods. The extracted samples are made of the first sentence and the matched sentences. The first sentence contains the true antecedent and the second the AZP.

Original text	لا بأس أن تُقالَ
pre-processed text	لا باس ان تقال

Table 3: An example of how we pre-process Arabic text. We normalize all 'alif' letters forms, and remove all diacritics.

lowed the pre-processing steps by [Althobaiti et al. \(2014\)](#) which led to a significant improvement in the performance of Arabic coreference resolution with an increase of 7% F1 score more than the baseline. Therefore, we follow their steps for pre-processing text which are:

- We normalize the various forms of the letter "alif" (أ، إ، آ) to the letter "ا".
- We remove all diacritic marks on words, such that each word has its undiacritized form.

We show an example of an original and pre-processed sentence from OntoNotes in Table 3.

5 Results

We use the systems from ([Aloraini and Poesio, 2020a](#)) and ([Aloraini and Poesio, 2020b](#)) as baselines for AZP identification and resolution respectively. To see the effect of the collected samples, we use their systems following the same settings. We evaluate the augmented data of each method with different settings, and see their performance on the precision, recall, and F1 scores.

AZP identification in Table 4, we train the baseline system with every data augmentation method and evaluate on the test portion. We see the results

vary. RSM and BT improve the scores with an increase of +0.4 and +0.1 compared to the baseline F1 score. ONP and CSA hurt the scores with a decrease of -0.3 and -0.7 compared to the baseline F1 score. When we train the system with ALL the augmented data, it outperforms the baseline with an increase of +0.3, but still less than when we use RSM by itself. RSM method provides the best performance and outperforms the baseline and all other settings scores.

AZP resolution we evaluate each augmented-data method and the results are shown in Table 5. Almost every data augmentation method increases the overall performance when it is trained with CoNLL-2012 except for CSA method. ONP, RSM, MCM, and BT increase the F1 score with +0.5, +1.6, +0.8, and +1.0 F1 scores respectively while CSA decreases the score by -0.2, compared with the baseline's F1 score. The best settings when we train ALL the augmented data with CoNLL-2012 which results an increase of +1.8.

Combinations we also try combining two methods together to see if they can outperform when we use every method separately. We show the AZP identification results in Table 6 and Table 7 for AZP resolution. As shown in the tables, some combinations can improve the scores; however, no single combination outperforms the highest scores in Tables 4 and 5.

5.1 Discussion

The augmented data helps to improve the results for both AZP identification and resolution. The improvements in scores for AZP resolution are more than AZP identification. To understand the reasons,

Training settings	Test Evaluation			
	P	R	F1	diff
CoNLL-2012 ^b	60.0	78.9	68.2	-
+ ONP	59.4	79.3	67.9	-0.3
+ RSM	59.8	80.4	68.6	+0.4
+ MCM	59.6	79.8	68.2	0
+ BT	59.6	80.2	68.3	+0.1
+ CSA	59.3	78.5	67.5	-0.7
CoNLL-2012 + ALL	59.8	80.2	68.5	+0.3

Table 4: AZP identification training settings with each data augmentation technique, and their results on the test set. ^b is the baseline scores (Aloraini and Poesio, 2020a). *diff* is the difference between a method’s F1 score and the baseline.

Training settings	Test Evaluation			
	P	R	F1	diff
CoNLL-2012 ^b	64.4	51.8	57.4	-
+ ONP	64.8	52.4	57.9	+0.5
+ RSM	65.6	53.7	59.0	+1.6
+ MCM	65.3	52.6	58.2	+0.8
+ BT	65.3	52.9	58.4	+1.0
+ CSA	64.4	51.6	57.2	-0.2
CoNLL-2012 + ALL	65.8	53.9	59.2	+1.8

Table 5: AZP resolution training settings with each data augmentation technique, and their results on the test set. ^b is the baseline scores (Aloraini and Poesio, 2020b). *diff* is the difference between a method’s F1 score and the baseline.

we investigated the Arabic portion of CoNLL-2012 and the outputs of the two systems. In CoNLL-2012, we found that many AZP samples to be duplicates but they refer to different types of mentions. The augmented data has more variance of AZP verbs; however, many were not present in the test part which might have led the model to associate these newly seen AZP verbs to cases that are not AZPs in the test set. We also examined the identified and resolved AZPs, and we found several cases to be verbs or mentions that are not subset of the training or the augmented data. Even though the data augmentation methods we showed helped to alleviate this problem, we still need more annotated data and explore more data augmentation methods.

Training settings	Test Evaluation			
	P	R	F1	diff
CoNLL-2012 ^b	60.0	78.9	68.2	-
+ ONP + RSM	59.6	79.7	68.2	0.0
+ ONP + MCM	59.5	79.7	68.1	-0.1
+ ONP + BT	59.8	79.8	68.3	+0.1
+ ONP + CSA	59.7	78.8	67.9	-0.3
+ RSM + MCM	59.8	79.0	68.1	-0.1
+ RSM + BT	59.8	79.8	68.4	+0.2
+ RSM + CSA	59.8	80.2	68.5	+0.3
+ MCM + BT	59.8	79.7	68.3	+0.1
+ MCM + CSA	59.7	79.7	68.2	0.0
+ BT + CSA	59.7	79.0	68.0	-0.2

Table 6: AZP identification training settings with different combinations of data augmentation methods, and their results on the test set.

Training settings	Test Evaluation			
	P	R	F1	diff
CoNLL-2012 ^b	64.4	51.8	57.4	-
+ ONP + RSM	65.3	53.7	58.9	+1.5
+ ONP + MCM	65.0	53.2	58.5	+1.1
+ ONP + BT	64.9	53.5	58.6	+1.1
+ ONP + CSA	64.6	52.7	58.0	+0.6
+ RSM + MCM	65.1	53.6	58.7	+1.4
+ RSM + BT	65.4	53.8	59.0	+1.6
+ MCM + BT	64.9	52.9	58.2	+0.8
+ MCM + CSA	65.1	53.4	58.6	+1.2
+ BT + CSA	64.6	52.4	57.8	+0.4

Table 7: AZP resolution training settings with different combinations of data augmentation methods, and their results on the test set.

6 Conclusion

We tested five methods to detect and generate AZPs. We suggested a simple and effective way to resolve AZPs in the automatically generated samples. We evaluated the augmented data on the AZP identification and resolution tasks for Arabic. Training the models with the augmented data achieved new state-of-the-art results for both AZP tasks.

Acknowledgements

We would like to thank the anonymous reviewers for their insightful comments and suggestions which helped to improve an early version of the paper.

References

- Bashir M. Alnajadat. 2017. Pro-drop in standard arabic. In *International Journal of English Linguistics 7.1*.
- Abdulrahman Aloraini and Massimo Poesio. 2020a. Anaphoric zero pronoun identification: A multilingual approach. In *Proceedings of the Third Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 22–32.
- Abdulrahman Aloraini and Massimo Poesio. 2020b. Cross-lingual zero pronoun resolution. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 90–98.
- Abdulrahman Aloraini, Juntao Yu, and Massimo Poesio. 2020. Neural coreference resolution for arabic. *arXiv preprint arXiv:2011.00286*.
- Maha Althobaiti, Udo Kruschwitz, and Massimo Poesio. 2014. Aranlp: A java-based library for the processing of arabic text.
- Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, Naama Tepper, and Naama Zwerdling. 2020. Do not have enough data? deep learning to the rescue! In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7383–7390.
- Jacob Andreas. 2019. Good-enough compositional data augmentation. *arXiv preprint arXiv:1904.09545*.
- Chinatsu Aone and Scott William Bennett. 1995. Evaluating automated and manual acquisition of anaphora resolution strategies. In *ACL '95 Proceedings of the 33rd annual meeting on Association for Computational Linguistics*, pages 122–129.
- Hitham M Abo Bakr, Khaled Shaalan, and Ibrahim Ziedan. 2009. A statistical method for detecting the arabic empty category. In *Proceedings of the Second International Conference on Arabic Language Resources and Tools, Cairo, Egypt*.
- Donna K. Byron, Whitney Gegg-Harrison, and Sun-Hee Lee. 2006. Resolving zero anaphors and pronouns in korean. In *Traitement Automatique des Langues 46.1*, pages 91–114.
- Tao Chang, Shaohe Lv, Xiaodong Wang, and Dong Wang. 2017. Zero pronoun identification in chinese language with deep neural networks. In *2017 2nd International Conference on Control, Automation and Artificial Intelligence (CAAI 2017)*. Atlantis Press.
- Chen Chen and Vincent Ng. 2013. Chinese zero pronoun resolution: Some recent advances. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1360–1365.
- Chen Chen and Vincent Ng. 2014. Chinese zero pronoun resolution: An unsupervised probabilistic model rivaling supervised resolvers. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.
- Chen Chen and Vincent Ng. 2015. Chinese zero pronoun resolution: A joint unsupervised discourse-aware model rivaling state-of-the-art resolvers. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 320–326.
- Chen Chen and Vincent Ng. 2016. Chinese zero pronoun resolution with deep neural network. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 778–788.
- Hannah Chen, Yangfeng Ji, and David Evans. 2020. Finding friends and flipping frenemies: Automatic paraphrase dataset augmentation using graph theory. *arXiv preprint arXiv:2011.01856*.
- Jiaao Chen, Derek Tam, Colin Raffel, Mohit Bansal, and Diyi Yang. 2021. An empirical survey of data augmentation for limited data learning in nlp. *arXiv preprint arXiv:2106.07499*.
- Susan Converse. 2006. Pronominal anaphora resolution in chinese. In *PhD Thesis, University of Pennsylvania*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- B. Di Eugenio. 1990. Centering theory and the italian pronominal system. In *Proc. of the 13th COLING, Helsinki, Finland*.
- Bosheng Ding, Linlin Liu, Lidong Bing, Canasai Kruengkrai, Thien Hai Nguyen, Shafiq Joty, Luo Si, and Chunyan Miao. 2020. Daga: Data augmentation with a generation approach for low-resource tagging tasks. *arXiv preprint arXiv:2011.01549*.
- Mushira Eid. 1983. On the communicative function of subject pronouns in arabic. In *Journal of Linguistics 19.2*, pages 287–303.
- Ibrahim Abu El-Khair. 2017. Abu el-khair corpus: A modern standard arabic corpus. *International Journal of Recent Trends in Engineering & Research (IJRTER) Volume, 2*.
- Steven Y Feng, Varun Gangal, Dongyeop Kang, Teruko Mitamura, and Eduard Hovy. 2020. Genau: Data augmentation for finetuning text generators. *arXiv preprint arXiv:2010.01794*.
- Steven Y Feng, Varun Gangal, Jason Wei, Sarath Chandar, Soroush Vosoughi, Teruko Mitamura, and Eduard Hovy. 2021. A survey of data augmentation approaches for nlp. *arXiv preprint arXiv:2105.03075*.
- Steven Y Feng, Aaron W Li, and Jesse Hoey. 2019. Keep calm and switch on! preserving sentiment and fluency in semantic text exchange. *arXiv preprint arXiv:1909.00088*.

- Antonio Ferrández and Jesús Peral. 2000. A computational approach to zero-pronouns in spanish. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, pages 166–172.
- Ryan Gabbard. 2010. Null element restoration. In *Ph.D Thesis, University of Pennsylvania*.
- Fei Gao, Jinhua Zhu, Lijun Wu, Yingce Xia, Tao Qin, Xueqi Cheng, Wengang Zhou, and Tie-Yan Liu. 2019. Soft contextual data augmentation for neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5539–5544.
- Madhav Gopal and Girish Nath Jha. 2017. Zero pronouns and their resolution in sanskrit texts. In *The International Symposium on Intelligent Systems Technologies and Application*, pages 255–267.
- Spence Green, Conal Sathi, and Christopher Manning. 2009. Np subject detection in verb-initial arabic clauses. In *Proceedings of the Third Workshop on Computational Approaches to Arabic Script-based Languages (CAASL3)*. Vol. 112.
- Diana Grigorova. 2013. An algorithm for zero pronoun resolution in bulgarian. In *Proceedings of the 14th International Conference on Computer Systems and Technologies*.
- Roman Grundkiewicz, Marcin Junczys-Dowmunt, and Kenneth Heafield. 2019. Neural grammatical error correction systems with unsupervised pre-training on synthetic data. In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 252–263.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. Colorless green recurrent networks dream hierarchically. *arXiv preprint arXiv:1803.11138*.
- Demi Guo, Yoon Kim, and Alexander M Rush. 2020. Sequence-level mixed sample data augmentation. *arXiv preprint arXiv:2011.09039*.
- Nizar Habash and Fatiha Sadat. 2006. Arabic pre-processing schemes for statistical machine translation. In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 49–52.
- Na-Rae Han. 2004. A korean null pronouns: Classification and annotation. In *Proceedings of the 2004 ACL Workshop on Discourse Annotation. Association for Computational Linguistics, 2004.*, pages 33–40.
- Masatsugu Hangyo, Daisuke Kawahara, and Sadao Kurohashi. 2013. Japanese zero reference resolution considering exophora and author/reader mentions. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 924–934.
- Jerry Hobbs. 1978. Resolving pronoun references. In *Lingua*, pages 311–338.
- Ryu Iida, Kentaro Inui, and Yuji Matsumoto. 2006. Exploiting syntactic patterns as clues in zero-anaphora resolution. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistic*, pages 625–632.
- Ryu Iida, Kentaro Inui, and Yuji Matsumoto. 2007. Zero-anaphora resolution by learning rich syntactic pattern features. In *ACM Transactions on Asian Language Information Processing*, 6(4).
- Ryu Iida and Massimo Poesio. 2011. A cross-lingual ilp solution to zero anaphora resolution. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 804–813.
- Ryu Iida, Kentaro Torisawa, Chikara Hashimoto, Jong-Hoon Oh, and Julien Kloetzer. 2015. Intra-sentential zero anaphora resolution using subject sharing recognition. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2179–2189.
- Hideki Isozaki and Tsutomu Hirao. 2003. Japanese zero pronoun resolution based on ranking rules and machine learning. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing.*, pages 184–191.
- Amit Jindal, Arijit Ghosh Chowdhury, Aniket Didolkar, Di Jin, Ramit Sawhney, and Rajiv Shah. 2020. Augmenting nlp models using latent feature interpolations. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6931–6936.
- Kushal Kaffle, Mohammed Yousefhusien, and Christopher Kanan. 2017. Data augmentation for visual question answering. In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 198–202.
- Megumi Kameyama. 1985. *Zero Anaphora: The case of Japanese*. Ph.D. thesis, Stanford University, Stanford, CA.
- Yeun-Bae Kim and Terumasa Ehara. 1995. Zero-subject resolution method based on probabilistic inference with evaluation function. In *Proceedings of the 3rd Natural Language Processing Pacific-Rim Symposium*, pages 721–727.
- YOUNG-JOO Kim. 2000. Subject/object drop in the acquisition of korean: A cross-linguistic comparison. In *Journal of East Asian Linguistics* 9.4, pages 325–351.
- Sosuke Kobayashi. 2018. Contextual augmentation: Data augmentation by words with paradigmatic relations. *arXiv preprint arXiv:1805.06201*.
- Fang Kong, Min Zhang, and Guodong Zhou. 2019. Chinese zero pronoun resolution: A chain-to-chain

- approach. *ACM Transactions on Asian and Low-Resource Language Information Processing (TAL-LIP)*, 19(1):1–21.
- Fang Kong and Guodong Zhou. 2010. A tree kernel-based unified framework for chinese zero anaphora resolution. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics*, pages 882–891.
- Ryuto Konno, Yuichiroh Matsubayashi, Shun Kiyono, Hiroki Ouchi, Ryo Takahashi, and Kentaro Inui. 2020. An empirical study of contextual data augmentation for japanese zero anaphora resolution. *arXiv preprint arXiv:2011.00948*.
- Ashutosh Kumar, Satwik Bhattamishra, Manik Bhandari, and Partha Talukdar. 2019. Submodular optimization-based diverse paraphrasing and its effectiveness in data augmentation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3609–3619.
- Sun-Hee Lee, Donna K Byron, and Seok Bae Jang. 2005. Why is zero marking important in korean? In *International Conference on Natural Language Processing*, pages 588–599. Springer.
- Charles N. Li and Sandra A. Thompson. 1979. Third person pronouns and zero anaphora in chinese discourse. In *Syntax and Semantics*, volume 12: Discourse and Syntax, pages 311–335. Academic Press.
- Ting Liu, Yiming Cui, Qingyu Yin, Weinan Zhang, Shijin Wang, and Guoping Hu. 2017. Generating and exploiting large-scale pseudo training data for zero pronoun resolution. In *arXiv preprint arXiv:1606.01603*.
- Samuel Louvan and Bernardo Magnini. 2020. Simple is better! lightweight data augmentation for low resource slot filling and intent classification. *arXiv preprint arXiv:2009.03695*.
- Christopher Manning and Hinrich Schütze. 1999. *Foundations of statistical natural language processing*. MIT press.
- Claudiu Mihăilă, Iustina Ilisei, , and Diana Inkpen. 2011. Zero pronominal anaphora resolution for the romanian language. In *Research Journal on Computer Science and Computer Engineering with Applications, POLIBITS*, 42.
- Ruslan Mitkov and Paul Schmidt. 1998. On the complexity of pronominal anaphora resolution in machine translation. *STUDIES IN FUNCTIONAL AND STRUCTURAL LINGUISTICS*, pages 207–222.
- Xuan-Phi Nguyen, Shafiq Joty, Wu Kui, and Ai Ti Aw. 2019. Data diversification: A simple strategy for neural machine translation. *arXiv preprint arXiv:1911.01986*.
- Ossama Obeid, Nasser Zalmout, Salam Khalifa, Dima Taji, Mai Oudah, Bashar Alhafni, Go Inoue, Fadhl Eryani, Alexander Erdmann, and Nizar Habash. 2020. Camel tools: An open source python toolkit for arabic natural language processing. In *Proceedings of the 12th language resources and evaluation conference*, pages 7022–7032.
- Luz Rello, Gabriela Ferraro, and Iria Gayo. 2012. A first approach to the automatic detection of zero subjects and impersonal constructions in portuguese. *Procesamiento del lenguaje natural*, 49:163–170.
- Luz Rello and Iustina Ilisei. 2009. A rule-based approach to the identification of spanish zero pronouns. In *Proceedings of the Student Research Workshop*, pages 60–65.
- Ryokan Ri, Toshiaki Nakazawa, and Yoshimasa Tsuruoka. 2021. Zero-pronoun data augmentation for japanese-to-english translation. *arXiv preprint arXiv:2107.00318*.
- Gözde Gül Şahin and Mark Steedman. 2019. Data augmentation via dependency tree morphing for low-resource languages. *arXiv preprint arXiv:1903.09460*.
- Ryohei Sasano, Daisuke Kawahara, and Sadao Kurohashi. 2008. A fully-lexicalized probabilistic model for japanese zero anaphora resolution. In *Proceedings of the 22nd International Conference on Computational Linguistics*, pages 769–776.
- Ryohei Sasano, Daisuke Kawahara, and Sadao Kurohashi. 2009. The effect of corpus size on case frame acquisition for discourse analysis. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 521–529.
- Ryohei Sasano and Sadao Kurohashi. 2011. discriminative approach to japanese zero anaphora resolution with large-scale lexicalized case frames. In *Proceedings of the 5th International Joint Conference on Natural Language Processing*, pages 758–766.
- Kazuhiro Seki, Atsushi Fujii, and Tetsuya Ishikawa. 2002. A probabilistic method for analyzing japanese anaphora integrating zero pronoun detection and resolution. In *Proceedings of the 19th International Conference on Computational Linguistics - Volume 1*, pages 1–7.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*.
- Jasdeep Singh, Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. 2019. Xlda: Cross-lingual data augmentation for natural language inference and question answering. *arXiv preprint arXiv:1905.11471*.

- Vaibhav Vaibhav, Sumeet Singh, Craig Stewart, and Graham Neubig. 2019. Improving robustness of machine translation with synthetic noise. *arXiv preprint arXiv:1902.09508*.
- Xinyi Wang, Hieu Pham, Zihang Dai, and Graham Neubig. 2018. Switchout: an efficient data augmentation algorithm for neural machine translation. *arXiv preprint arXiv:1808.07512*.
- Jason Wei and Kai Zou. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196*.
- Lydia White. 1985. The “pro-drop” parameter in adult second language acquisition. *Language learning*, 35(1):47–61.
- Ziang Xie, Guillaume Genthial, Stanley Xie, Andrew Y Ng, and Dan Jurafsky. 2018. Noising and denoising natural language: Diverse backtranslation for grammar correction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 619–628.
- Souta Yamashiro, Hitoshi Nishikawa, and Takenobu Tokunaga. 2018. Neural japanese zero anaphora resolution using smoothed large-scale case frames with word embedding. In *Proceedings of the 32nd Pacific Asia Conference on Language, Information and Computation*.
- Wei Yang, Yuqing Xie, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019. Data augmentation for bert fine-tuning in open-domain question answering. *arXiv preprint arXiv:1904.06652*.
- Ching-Long Yeh and Yi-Chun Chen. 2006. Zero anaphora resolution in chinese with shallow parsing. In *Journal of Chinese Language and Computing 17 (1)*, pages 41–56.
- Qingyu Yin, Yu Zhang, Weinan Zhang, and Ting Liu. 2016. A deep neural network for chinese zero pronoun resolution. In *arXiv preprint arXiv:1604.05800*.
- Qingyu Yin, Yu Zhang, Weinan Zhang, and Ting Liu. 2017. Chinese zero pronoun resolution with deep memory network. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1309–1318.
- Qingyu Yin, Yu Zhang, Weinan Zhang, Ting Liu, and William Yang Wang. 2018. Zero pronoun resolution with attention-based neural network. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 13–23.
- Katsumasa Yoshikawa, Masayuki Asahara, and Yuji Matsumoto. 2011. Jointly extracting japanese predicate-argument relation with markov logic. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 1125–1133.
- Kei Yoshimoto. 1988. Identifying zero pronouns in japanese dialogue. In *Coling Budapest 1988 Volume 2: International Conference on Computational Linguistics*.
- Koichiro Yoshino, Shinsuke Mori, and Tatsuya Kawahara. 2013. Predicate argument structure analysis using partially annotated corpora. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 957–961.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28:649–657.
- Yi Zhang, Tao Ge, Furu Wei, Ming Zhou, and Xu Sun. 2019. Sequence-to-sequence pre-training with data augmentation for sentence rewriting. *arXiv preprint arXiv:1909.06002*.
- Shanheng Zhao and Hwee Tou Ng. 2007. Identification and resolution of chinese zero pronouns: A machine learning approach. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 541–550.