



This is a repository copy of *Free the plural: unrestricted split-antecedent anaphora resolution*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/190603/>

Version: Published Version

Proceedings Paper:

Yu, J., Moosavi, N.S. orcid.org/0000-0002-8332-307X, Paun, S. et al. (1 more author) (2020) Free the plural: unrestricted split-antecedent anaphora resolution. In: Scott, D., Bel, N. and Zong, C., (eds.) Proceedings of the 28th International Conference on Computational Linguistics. COLING'2020: 28th International Conference on Computational Linguistics, 08-13 Dec 2020, Barcelona, Spain (Online). International Committee on Computational Linguistics , pp. 6113-6125. ISBN 9781952148279

<https://doi.org/10.18653/v1/2020.coling-main.538>

© 2020 The Author(s). Available under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>).

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Free the Plural: Unrestricted Split-Antecedent Anaphora Resolution

Juntao Yu¹, Nafise Sadat Moosavi², Silviu Paun¹, Massimo Poesio¹

¹Queen Mary University of London

²UKP Lab, Technische Universität Darmstadt

¹{juntao.yu, s.paun, m.poesio}@qmul.ac.uk

²moosavi@ukp.informatik.tu-darmstadt.de

Abstract

Now that the performance of coreference resolvers on the simpler forms of anaphoric reference has greatly improved, more attention is devoted to more complex aspects of anaphora. One limitation of virtually all coreference resolution models is the focus on single-antecedent anaphors. Plural anaphors with multiple antecedents—so-called split-antecedent anaphors (as in *John met Mary. They went to the movies*)—have not been widely studied, because they are not annotated in ONTONOTES and are relatively infrequent in other corpora. In this paper, we introduce the first model for unrestricted resolution of split-antecedent anaphors. We start with a strong baseline enhanced by BERT embeddings, and show that we can substantially improve its performance by addressing the sparsity issue. To do this, we experiment with auxiliary corpora where split-antecedent anaphors were annotated by the crowd, and with transfer learning models using **element-of** bridging references and single-antecedent coreference as auxiliary tasks. Evaluation on the gold annotated ARRAU corpus shows that the out best model uses a combination of three auxiliary corpora achieved F1 scores of 70% and 43.6% when evaluated in a lenient and strict setting, respectively, i.e., 11 and 21 percentage points gain when compared with our baseline.¹

1 Introduction

(Identity) anaphora resolution (coreference) is the linguistic task of linking nominal expressions (mentions) to entities in the discourse, so that mentions representing the same entity are grouped together in a ‘coreference chain’ (Poesio et al., 2016). As the performance of coreference models has substantially improved (Clark and Manning, 2015; Clark and Manning, 2016; Lee et al., 2017; Lee et al., 2018; Kantor and Globerson, 2019) in recent years, more attention is being devoted to more complex aspects of anaphoric reference— from the pronouns that require commonsense knowledge for their resolution studied in the Winograd Schema Challenge (Rahman and Ng, 2012; Peng et al., 2015) to pronouns that cannot be resolved purely on the basis of gender (Webster et al., 2018). Another limitation of state of the art systems is the assumption that anaphors can only have one antecedent. Plural anaphors with multiple antecedents (**split antecedent anaphors**) are not widely studied— in fact, such anaphors are not annotated in the most widely used coreference corpus, ONTONOTES (Pradhan et al., 2012). In ONTONOTES we find annotated cases of plural reference to plural antecedents, as in (1), or cases in which singular antecedents are conjoined so that a mention can be introduced for the conjunction, as in (2). However, it is also possible to refer plurally to antecedents introduced by separate noun phrases, as in (3) or (4) (Eschenbach et al., 1989; Kamp and Reyle, 1993); such cases are not annotated in ONTONOTES.

- (1) *The Joneses_i* went to the park. They_i had a good time.
- (2) *John and Mary_i* went to the park. They_i had a good time.
- (3) *John_i* met *Mary_j* in the park. They_{i,j} had a good chat .
- (4) John likes *green_i*, Mary likes *blue_j*, but Tom likes both colours_{i,j}.

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

¹The code is available at <https://github.com/juntaoy/dali-plural>

Early research on split antecedent anaphors mostly focused on the constraints on the construction of complex entities from singular entities. There are two recent studies of split antecedent anaphora, both involving the creation of a new dataset, and focusing on a subset of the problem. Vala et al. (2016) proposed a model focused on the resolution of the pronominal plural mentions *they* and *them* in a newly created corpus of fiction where most of the antecedents are from a fixed list of characters in the novel. Zhou and Choi (2018) proposed a model that addresses a wider range of anaphoric references to split antecedents, but limited only to references to main characters (mainly pronominal mentions) in a corpus of transcripts of the sitcom *Friends*.

In this paper, we introduce the first model targeting the whole range of split-antecedent anaphora. We evaluate our system on the hand-annotated ARRAU corpus (Poesio and Artstein, 2008; Uryupina et al., 2020), which covers a range of anaphoric relations going from identity relations (including split antecedent anaphora), bridging reference, and discourse deixis, as well as different genres going from news to task-oriented dialogue. Since the task is complex, we focus on establishing links between (gold) plural anaphors and their split-antecedents, leaving the detection of split-antecedent anaphors for future work. We follow Vala et al. (2016) and evaluate on the setting that assumes the gold split-antecedent anaphors and the gold mentions are provided.

Our baseline system is a simplified version of the state of the art coreference system (Lee et al., 2018; Kantor and Globerson, 2019) enhanced by BERT embeddings (Devlin et al., 2019). The key issue we tackled is that compared to single-antecedent anaphors, the number of split-antecedent anaphors is rather small: in total, only 697 split-antecedent anaphors were annotated in the ARRAU corpus, about 2% of all anaphoric references. To tackle this challenge of limited training data, we experimented with different ways of **using auxiliary corpora** to improve performance. We evaluated four different augmentation settings. Two of these involve using additional examples of split antecedent anaphora recoverable from the crowdsourced *Phrase Detectives* corpus (PD) (Poesio et al., 2019), a corpus of anaphoric annotations, including split-antecedent anaphors, collected using the *Phrase Detectives* game.² The corpus includes both raw annotations and silver labels aggregated using the Mention Pair Annotations model (Paun et al., 2018). For our first setting (PD-SILVER), we used the silver labels to identify split-antecedent anaphors in the corpus, and use them as an auxiliary corpus. For the second setting (PD-CROWD), we use the same data as PD-SILVER, but instead of using silver labels we collect all ‘raw’ split-antecedent annotations, and use majority voting to choose the labels when there are different annotations for the same anaphor. The third and fourth approaches to augmentation involve a form of transfer learning, using annotations of different but related phenomena to help learning split antecedent resolution. In our third setting (ELEMENT-OF), we used as auxiliary training data examples of a type of bridging reference, **element-of**, that is related to our task and is annotated in the ARRAU corpus.³ Finally, in our last setting (SINGLE-COREF), we used as auxiliary training data the examples of single-antecedent coreference annotated in the ARRAU corpus.

We also evaluated three different training strategies to leverage the auxiliary data. The first strategy (CONCAT) involves randomly selecting a training document from the main corpus and the auxiliary corpus in turns, with a fixed probability of 0.5 to train with the main corpus. For the second strategy (PRE-TRAIN), we first pre-train our model on the auxiliary corpus, and fine-tune it on the main corpus. Our last strategy (ANNEALING) is inspired by the teacher annealing method used in Clark et al. (2019). We train our models with both the main corpus and the auxiliary corpus as in the CONCAT strategy, but linearly increase the usage of documents from the main corpus. In this way, the training progressively switches from the auxiliary corpus to the main corpus.

The evaluation on the ARRAU corpus shows that all four auxiliary datasets combined with all three training strategy substantially improve the performance of the baseline model. The evaluation on mixed use of our auxiliary corpora results in further improvements on performance. Our best model, trained with three auxiliary corpora (PD-CROWD, ELEMENT-OF, SINGLE-COREF), outperforms our strong baseline by 11.4 and 20.9 percentage points when evaluated in a lenient and strict setting, respectively. The

²<https://www.phrasedetectives.org/>

³The relation between the split-antecedent anaphors and the individual antecedents is the inverse of **element-of**.

final model has an F1 score of 70% when partial credit is granted in the evaluation, and correctly resolves all antecedents for 43.6% of the split-antecedent anaphors (the strict evaluation). To the best of our knowledge, this is the first reported result on unrestricted split-antecedent anaphora resolution.

2 Background

2.1 Single-Antecedent Coreference Resolution

Single-antecedent coreference resolution has been extensively studied. In the pre-neural period, both rule-based (Lee et al., 2013) and statistical models (Soon et al., 2001; Björkelund and Kuhn, 2014; Clark and Manning, 2015) were developed to resolve single-antecedent coreference resolution. Recently, Wiseman et al. (2015; Wiseman et al. (2016) first introduced a neural network-based approach to solving coreference in a non-linear way. Clark and Manning (2016) integrated reinforcement learning to let the model, optimising directly on the B³ scores. Lee et al. (2017) proposed a neural joint approach for mention detection and coreference resolution. Their model does not rely on parse trees; instead, the system learns to detect mentions by exploring the outputs of a BILSTM. After the introduction of contextual word embeddings such as ELMO (Peters et al., 2018) and BERT (Devlin et al., 2019), the Lee et al. (2017) system has been greatly improved by those embeddings (Lee et al., 2018; Kantor and Globerson, 2019; Joshi et al., 2019; Joshi et al., 2020) to achieve SoTA results. But none of those SoTA systems can resolve split-antecedent anaphors.

2.2 Split-Antecedent Anaphora

There is substantial research on resolving split-antecedent anaphors both in linguistics (Kamp and Reyle, 1993) and psychology (Murphy, 1984; Sanford and Lockhart, 1990; Kaup et al., 2002; Patson, 2014), but only a few early computational models (Eschenbach et al., 1989). This work was primarily concerned with explaining preferences and restrictions on split antecedent anaphora: for example, in (5a) *they* can refer to Michael, Peter and Maria, or to Peter and Maria, but not to Michael and Maria; or in (5b) there seems to be a preference for *they* to refer to Peter and Maria (Kaup et al., 2002). Proposals differed, e.g., on whether complex reference objects are created immediately or after encountering the anaphor.

- (5) a. Michael met Peter and Maria in the pub. They had a great time.
b. Michael watched Peter and Maria in the pub. They had a great time.

Only two recent computational treatments of this type of anaphor exist (Vala et al., 2016; Zhou and Choi, 2018); they are discussed in Section 6.

2.3 Approaches for Under-Resourced Tasks

Much research has focused on resolving under-resourced tasks via semi-supervised learning (Pekar et al., 2014; Yu and Bohnet, 2015; Kocijan et al., 2019; Hou, 2020) and shared representation based transfer learning (Yang et al., 2017; Cotterell and Duh, 2017; Zhou et al., 2019).

Semi-supervised learning Semi-supervised methods use large unlabelled/automatically labelled data to enhance performance for under-resourced domains/languages. Early research focused on generating additional training data from automatically annotated text. Pekar et al. (2014), Yu et al. (2015) used co-training/self-training for dependency parsing to leverage models trained on rich-resourced domain for under-resourced domains. Yu and Bohnet (2015) applied a confidence-based self-training approach to enhance parsing performance for 9 languages with parse trees automatically annotated by models trained on small initial training data.

Recently, another line of research focused on creating synthetic training data from unlabelled/automatically labelled data using some heuristic patterns. Kocijan et al. (2019) used Wikipedia to create WikiCREM, a large pronoun resolution dataset, using heuristic rules based on the occurrence of personal names in sentences. The evaluation of their system on Winograd Schema corpora shows that models pre-trained on the WikiCREM consistently outperform models that do not use it. Hou (2020) applied a similar approach to the antecedent selection task for bridging references. Hou created an artificial bridging corpus using the prepositional and possessive structures in the automatically parsed Gigaword

corpus. Models pre-trained with the artificial corpus achieved substantial gains over baselines. Our PD-SILVER and PD-CROWD settings are close to this approach, but instead of using automatically annotated data, we use crowdsourced annotations from the PD corpus. Both our corpora and the synthetic corpora created by Hou (2020) contain a degree of noise compared with gold annotated corpora.

Shared Representation-Based Transfer Learning Shared representation-based transfer learning focuses on exploiting auxiliary tasks for which large annotated data exist to help in under-resourced tasks/domains/languages. It is similar to multi-task learning, but only focuses on enhancing the performance of the under-resourced task. Yang et al. (2017) applied transfer learning to sequence labelling tasks; the deep hierarchical recurrent neural network used in their work is fully/partially shared between the source and the target tasks. They demonstrated that SoTA performance can be achieved by using models trained on multi-tasks. Cotterell and Duh (2017) trained a neural NER system on a combination of high-/low-resource languages to improve NER for the low-resource languages. In their work, character-based embeddings are shared across the languages. Recently, Zhou et al. (2019) introduced a multi-task network together with adversarial learning for under-resourced NER. The evaluation on both cross-language and cross-domain settings shows that partially sharing the BILSTM works better for cross-language transfer, while for cross-domain setting, the system performs better when the LSTM layers are fully shared. Our third and fourth settings (ELEMENT-OF and SINGLE-COREF) can be viewed as shared representation-based transfer learning where we use bridging resolution and single-antecedent coreference resolution as our auxiliary tasks to aid our split-antecedent anaphora resolution.

3 Methods

3.1 The Baseline System

Our baseline is a simplified version of the SoTA coreference architecture by Lee et al. (2018), further developed by Kantor and Globerson (2019). In this model, mention detection and coreference are carried out jointly, but here we only use the coreference part, since we evaluate our model with gold mentions.

Our baseline system first creates representations for mentions using the output of a BILSTM. The BILSTM takes as input the concatenated embeddings of both word and character levels. For word embeddings, GloVe (Pennington et al., 2014) and BERT (Devlin et al., 2019) embeddings are used. Character embeddings are learned from a convolution neural network (CNN) during training. The tokens are represented by concatenating outputs from the forward and the backward LSTMs. The token representations $(x_t)_{t=1}^T$ are used together with head representations (h_i) to represent mentions (M_i) . The h_i of a mention is obtained by applying attention over its token representations $(\{x_{b_i}, \dots, x_{e_i}\})$, where b_i and e_i are the indices of the start and the end of the mention, respectively. Formally, we compute h_i, M_i as follows:

$$\begin{aligned}\alpha_t &= \text{FFNN}_\alpha(x_t) \\ a_{i,t} &= \frac{\exp(\alpha_t)}{\sum_{k=b_i}^{e_i} \exp(\alpha_k)} \\ h_i &= \sum_{t=b_i}^{e_i} a_{i,t} \cdot x_t \\ M_i &= [x_{b_i}, x_{e_i}, h_i, \phi(i)]\end{aligned}$$

where $\phi(i)$ is the mention width feature embeddings. Next, we pair the mentions with candidate antecedents to create a pair representation $(P_{(i,j)})$:

$$P_{(i,j)} = [M_i, M_j, M_i \circ M_j, \phi(i, j)]$$

where M_i, M_j is the representation of the antecedent and anaphor, respectively, \circ denotes element-wise product, and $\phi(i, j)$ is the distance feature between a mention pair. To make the model computationally tractable, we consider for each mention a maximum 250 candidate antecedents as we observed from the ARRAU corpus, most of the antecedents can be retrieved within the 250 candidates window size.

The next step is to compute the pairwise score ($s(i, j)$). Following Lee et al. (2018), we add an artificial antecedent ϵ to deal with cases of non-split-antecedent anaphor mentions or cases when the antecedent does not appear in the candidate list during the training. We do not use ϵ during test time as we use gold split-antecedent anaphors. We compute $s(i, j)$ as follows:

$$r(i, j) = \text{FFNN}(P_{(i,j)})$$

$$s(i, j) = \frac{1}{1 + e^{-r(i,j)}}$$

At test time, the system will generate two to five antecedents according to their $s(i, j)$ scores. The upper threshold is based on the observation that the vast majority of split-antecedent anaphors in ARRAU have no more than 5 antecedents. To generate the antecedents, we first rank the candidates by their $s(i, j)$ scores in descending order. We add up to 5 top candidates that have a $s(i, j)$ score above 0.5.⁴ If less than two candidates were selected, we add top two candidates to the predictions regardless of their scores.

3.2 Auxiliary Corpora

Since the number of examples of split-antecedent anaphora in ARRAU is small, we deployed four auxiliary corpora created from either the crowd annotated *Phrase Detectives* (PD) corpus or the gold annotated ARRAU corpus to improve the performance of the system.

The PD corpus was created using the *Phrase Detectives* game, whose players are asked to find the antecedent/split-antecedents closest to the mention in question (Poesio et al., 2019). The corpus comes with all raw annotations and silver labels aggregated using the Mention-Pair Annotation model (Paun et al., 2018). We created our first two auxiliary corpora from the latest version of the PD corpus⁵ by using different aggregation methods. The ARRAU corpus consists of texts from four very distinct domains: news (the RST subcorpus), dialogue (the TRAINS subcorpus), fiction (the PEAR stories), and medical / art history (the GNOME subcorpus). Its annotation scheme covers the annotation of referring (including singletons) and non-referring expressions; coreference relations including split antecedent plurals and generic references; and non-coreferential anaphoric relations including discourse deixis and bridging references. We create the other two auxiliary corpora from the ARRAU corpus. The rest of the subsection describes our auxiliary corpora in detail.⁶

Silver Labels (PD-SILVER) For our first auxiliary corpus, we simply added to our training data the split-antecedent anaphora examples from the PD corpus. We used the silver labels that come with the corpus and extracted 507 split-antecedent anaphors (see Table 1). This nearly doubled the size of our training data. We assessed the quality of the silver labels by comparing it against the gold annotated subset of the PD corpus;⁷ the silver labels have a relatively good quality (62.9% F1), recalling 68.8% of the gold split-antecedent anaphoric links and with a precision of 57.9%.

Raw Crowd Annotations (PD-CROWD) The second auxiliary corpus was created by extracting all split-antecedent examples from the raw annotations in PD to maximise recall. After extracting all split-antecedent annotations, we used majority voting as our aggregation method when players did not agree on split-antecedent annotations. In this way, we extracted 47.7k split-antecedent annotations associated with 6.2k mentions (Table 1). The quality of this extraction method was evaluated on the gold portion of the PD corpus as well; the resulting dataset has a recall of 91.7%, which fulfils the goal of this setting. As expected, the corpus is noisy, with a precision of 11.1% and an F1 of 19.7%. We manually checked the false-positive examples, finding they are mainly due to three types of mistakes: single-antecedent coreference (the coreference chains were annotated as the split-antecedent), bridging reference (not required to be annotated), and other annotation mistakes. The first two types of mistakes are not harmful to our task as our third and fourth auxiliary corpora are created using those types of relations.

⁴We use the gold single-antecedent clusters to ensure the selected antecedents belong to distinct gold cluster.

⁵<https://github.com/dali-ambiguity/Phrase-Detectives-Corpus-2.1.4>

⁶The ParCorFull corpus (Lapshinova-Koltunski et al., 2018) also includes split-antecedent annotations, but the number of split-antecedent examples are too small to be used in our experiments.

⁷A subset of the PD corpus comes with additional gold labels annotated by experts.

Corpus	Anaphors Type	Data Quality	Num of docs	Num of Anaphors
ARRAU TRAIN/DEV/TEST	Split-antecedent anaphors	Gold	211 / 30 / 60	507 / 80 / 110
PD-SILVER	Split-antecedent anaphors	Silver	165	507
PD-CROWD	Split-antecedent anaphors	Noisy	467	6262
ELEMENT-OF	Bridging anaphors	Gold	213	1059
SINGLE-COREF	Single-antecedent anaphors	Gold	462	30372

Table 1: Statistics about the corpora used in our experiments.

Element-of Bridging References (ELEMENT-OF) ARRAU is also annotated with bridging references, and one of the bridging relations covered by the annotation, **element-of** (and its inverse) are very closely related to the task of resolving split-antecedent plurals. **Element-of** is the relation between a new singular entity and a plural entity introduced in the discourse, as in (6a), or between a previously introduced singular entity and a new plural entity, as in (6b).

- (6) a. There are *two supermarkets* in our village, but one is very small. (**element-of**)
b. Yet *another small bookshop* just opened in our village. Our independent bookshops are our main attraction. (**element-of-inverse**)

Since the proposed system uses a pairwise approach, the relations between split-antecedent anaphors and their antecedents are established by multi-links between anaphors and individual antecedents. These are **element-of** relations, but differ from the bridging case in two respects. First, the plural relations are an inverse version of the element-of relations where the antecedent is an element of the anaphor. Second, split-antecedent coreference *is* coreference: the union of all antecedents has the same denotation as the anaphor, unlike in bridging. Nevertheless, the **element-of** bridging relation is close enough to be potentially useful for our task. We therefore created a third auxiliary corpus by extracting **element-of** bridging relations from ARRAU. In total we extracted 1059 training examples (see Table 1).

Single-antecedent anaphors (SINGLE-COREF) Our last auxiliary corpus using single-antecedent anaphors. The main reason for using single-antecedent anaphors as supporting dataset is that single-antecedent anaphors are very common: e.g., in ARRAU TRAIN we only have 500 split-antecedent anaphors, but 30k single-antecedent anaphors (see Table 1). This gives us a much larger corpus than all other auxiliary corpora proposed earlier. Using a large auxiliary corpus allows our system to learn a better mention and pairwise representations that might be beneficial for our under-resourced task.

3.3 Training Strategies

Training with multiple corpora is challenging, especially when the auxiliary corpus is noisy. In this paper, we evaluate our system with three different training strategies to maximise the performance on split-antecedent anaphora resolution.

Concatenation (CONCAT) The first and simplest strategy is to use the auxiliary corpus as additional training data by concatenating it with the main corpus. We configured training to train on documents from the main and auxiliary corpus in turn, with 50% of the time on the main corpus. By doing so, we make sure the system will not overfit the auxiliary corpus.

Pre-training (PRE-TRAIN) Our second strategy was to first pre-train the system on the auxiliary corpus, and then fine-tuning the model on the main corpus to fit our task. Such a strategy works well when the auxiliary corpus is noisy as the fine-tuning step will only be trained on the gold annotations.

Corpus Annealing (ANNEALING) Our last strategy was inspired by Clark et al. (2019)’s teacher annealing proposal. Clark et al. (2019) use teacher annealing to enable smoother learning. The multi-task model initially learns from the predictions of the single-task model, but training gradually switches to gold labels by a weighted loss function. In this paper, we configured our system to initially learn from the auxiliary corpus, and using a linearly decreasing ratio of training with the auxiliary corpus. Instead of using a weighted loss as done by Clark et al. (2019), we used the ratio to control the source of our training documents (main or auxiliary). By doing so, the learning process smoothly switches from the auxiliary corpus to the main corpus, training 100% on the main corpus when the end of training is reached.

3.4 Learning

Following Lee et al. (2018), we optimise our model on the marginal log-likelihood of all correct antecedents. We consider an antecedent correct if it is from the same gold single-antecedent coreference cluster $\text{GOLD}(i)$ as the gold antecedent lists. We also use the gold single-antecedent clusters to extend the split-antecedent anaphor list during training: i.e., mentions in the same single-antecedent cluster of a split-antecedent anaphor are considered as split-antecedent anaphors. The extension boosts the number of split-antecedent anaphors in the training data by 79% to 908. We compute the losses as follows:

$$\log \prod_{j=1}^N \sum_{\hat{y} \in Y(j) \cap \text{GOLD}(i)} r(\hat{y}, j)$$

in case mention i is not a split-antecedent anaphor or $Y(j)$ (the candidate antecedents) does not contain mentions from $\text{GOLD}(i)$, we set $\text{GOLD}(i) = \{\epsilon\}$.

4 Experiments

Datasets We evaluated our models on the ARRAU corpus (Poesio and Artstein, 2008; Uryupina et al., 2020) as this is the only gold annotated corpus with split-antecedent anaphors annotated.⁸ The corpus also contains annotations of the bridging references and single-antecedent coreference relations used for the auxiliary datasets. We used all four subcorpora of ARRAU: RST (news), TRAINS (dialogue), PEAR (fiction) and GNOME (medical and art history). 301 out of 552 total documents contain split-antecedent anaphors. We use the 1-7th, 8th, 9-10th of every 10 documents as our train, dev and test dataset, respectively (see Table 1 for more detail).

In addition, we used the *Phrase Detectives* corpus to create auxiliary datasets. The PD corpus contains 542 documents from two main domains, Wikipedia and fiction. 165 documents contain split-antecedent anaphors according to the silver labels in the corpus; we use those documents as our PD-SILVER corpus. Our PD-CROWD auxiliary corpus consists of the 467 documents which contain split antecedents when aggregated as described in Section 3.2. The ELEMENT-OF corpus has 213 documents contain element-of bridging relations from the non-dev/test portion of the ARRAU corpus. The SINGLE-COREF corpus is formed by 462 non-dev/test documents of the ARRAU corpus. Table 1 shows statistics about our corpora.

Evaluation metrics Following Vala et al. (2016) we report **lenient** F1 scores that give partial credit when only some of the split individual antecedents of a plural are found, and consider an antecedent correct as long as it belongs to the correct gold single-antecedent cluster. We further report **strict** scores that require all antecedents of split-antecedent anaphors be correctly resolved for the final evaluation.

Hyperparameters We used the default settings from Lee et al. (2018), replacing their ELMO settings with the BERT settings from Kantor and Globerson (2019). We trained the models (including pre-training models) for 200k steps.

5 Results and Discussions

5.1 Training Strategy Selection

We first applied all three training strategies to our auxiliary corpora to find the best training strategy for each auxiliary corpus. We used lenient F1 scores on the development set to select the strategy most suitable for each individual corpus. As illustrated in Table 2, our baseline model trained only on ARRAU TRAIN already achieves a reasonably good F1 score for this task (58.2%). Starting with a strong baseline, our system enhanced by the auxiliary corpora achieved substantial improvements of up to 11.3%.

Among the training strategies, PD-SILVER works best when using the CONCAT method. This makes sense as PD-SILVER contains split antecedent examples annotated using the same annotation scheme as ARRAU. The PD-CROWD corpus is much noisier, but despite containing a large amount of false positives, it achieves better F1 scores than PD-SILVER. This confirms our hypothesis that a higher recall of split-antecedent examples is important, and the false positive examples –mainly single-antecedent anaphors

⁸As far as we know the corpus used by Vala et al. (2016) is not publically available, and the corpus used by Zhou and Choi (2018) only covers anaphoric references to a limited range of antecedents.

Training Strategy	BASELINE	PD-SILVER	PD-CROWD	ELEMENT-OF	SINGLE-COREF
CONCAT	58.2	59.8	61.2	59.2	67.6
PRE-TRAIN	58.2	59.0	62.9	64.3	66.5
ANNEALING	58.2	59.0	62.6	61.1	69.5

Table 2: Training strategy selection on the development set with lenient F1 scores.

	Lenient			Strict
	R	P	F1	Accuracy
RECENT-2	19.6	21.8	20.6	3.6
RECENT-3	31.8	23.6	27.1	0.9
RECENT-4	40.4	22.6	28.9	0.0
RECENT-5	45.7	20.4	28.2	0.0
RANDOM	24.9	11.4	15.7	0.0
NEURAL BASELINE	60.8	56.4	58.6	22.7
PD-SILVER	61.6	61.9	61.8	30.9
PD-CROWD	68.2	63.5	65.7	31.8
ELEMENT-OF	64.5	65.0	64.8	34.5
SINGLE-COREF	68.6	70.6	69.6	42.7
PD-CROWD + SINGLE-COREF	68.2	69.6	68.9	40.9
ELEMENT-OF + SINGLE-COREF	69.4	67.5	68.4	39.1
PD-CROWD + ELEMENT-OF + SINGLE-COREF	72.2	67.8	70.0	43.6

Table 3: Comparing our models with the baselines on the test set.

and bridging relations—do not harm the results. Both PRE-TRAIN and ANNEALING are suitable strategies for the PD-CROWD corpus, with the former slightly better. The ELEMENT-OF corpus works best in a PRE-TRAIN setting, with a large improvement of 6.2% when compared to the baseline even though the corpus only contains a small number of examples (1k). The large improvement confirmed our hypothesis that **element-of** bridging relations are closely related to split-antecedent relations. Finally, the SINGLE-COREF corpus achieved the best scores with all three training strategy, but the largest improvement of 11.3% is achieved by training with ANNEALING method. As the SINGLE-COREF corpus has a substantially larger number of examples when compared to all other auxiliary corpus used in this paper, it seems likely that this might be an important reason for its usefulness. Overall, our auxiliary corpora and training strategies showed their merit for enhancing the performance on split-antecedent anaphora resolution; we will further discuss this in later sections.

5.2 Comparison with the Baselines

We then evaluated our models, each trained using the best training strategy for that corpus, on the test set. Since our paper reports the first result on split-antecedent anaphora resolution on ARRAU, we compare our system with various baselines. Following Vala et al. (2016), we created naive baselines: RECENT-M and RANDOM. RECENT-M assigns to an anaphor the m antecedents (from distinct single-antecedent clusters) that are closest. RANDOM assigns all candidate antecedents random probabilities; the antecedents are selected using the same method as the one used in our trained models.

Table 3 shows the results on the test set. The naive baselines achieved a maximum lenient F1 score of 28.9% when using the 4 most recent antecedents. When using strict evaluation, most naive baselines perform really badly. This poor performance of the naive baselines confirms the difficulty of our task. Our NEURAL BASELINE trained solely on ARRAU TRAIN achieved a lenient F1 of 58.6%, more than double the best result of the naive baselines. With strict evaluation, the same model achieved 22.7%—6 times the best score of naive baselines, but still low. Using auxiliary corpora improved the performance of the neural model by a minimum of 3.2 and 8.2 p.p. when evaluated in a lenient and strict setting, respectively. But our best model, trained using SINGLE-COREF auxiliary information, achieved a lenient

	Count	NEURAL BASELINE		BEST	
		Lenient	Strict	Lenient	Strict
2	157	60.1	33.1	69.5	44.0
3+	33	52.1	6.1	65.3	21.2

(a) Scores for anaphors with different number of antecedents.

	1k	6k	10k	20k	All (30k)
Lenient	60.1	63.3	65.9	67.0	69.5
Strict	28.8	32.5	40.0	46.3	46.3

(b) Scores of models trained with reduced size of the SINGLE-COREF corpus.

Table 4: Analysis of our best model.

F1 of 69.6% and a strict accuracy of 42.7%, which is 11% and 20% better than our NEURAL BASELINE.

We further evaluated using combinations of auxiliary corpora—i.e., using the PD-CROWD, ELEMENT-OF, and SINGLE-COREF corpora, and using either the PRE-TRAIN or ANNEALING strategies—e.g., using PD-CROWD for PRE-TRAIN, then fine-tune the model with SINGLE-COREF corpus using ANNEALING strategy. In total, we evaluated three combinations (see Table 3). The best result, achieved by combining all three auxiliary corpora, was 0.4 p.p. and 0.9 p.p. better than the results achieved by using SINGLE-COREF alone in a lenient and strict evaluation, respectively.

5.3 Analysis

Number of Antecedents We compared our best model with the neural baseline, using both lenient and strict scores and also considering the number of split-antecedents. For this analysis, we evaluated both models on the concatenation of test and development set to collect more examples.

As shown in Table 4a, 82.6% of the anaphors in the dataset have two antecedents; the rest have three or more antecedents. With anaphors that have two antecedents, our best model achieved improvements of 9.5% (lenient) and 10.9% (strict), respectively. With anaphors that have more antecedents, our best model achieved even larger improvements for both lenient (13.2%) and strict (15.1%) scores. Overall, our best model outperforms the baseline by large margins in all evaluations. System prediction examples for both systems can be find in Table 5.

Size of the Auxiliary Corpus Our SINGLE-COREF auxiliary corpus achieved much larger improvements than all other corpora evaluated in this paper. A simple explanation would be that this is because the SINGLE-COREF corpus is substantially larger. So, to understand the impact of the auxiliary corpus size on our task, we further trained our model with auxiliary corpora of different size. The examples are randomly selected from our SINGLE-COREF corpus. Table 4b shows our results on the development set. When using 1k examples from the SINGLE-COREF corpus, the lenient F1 is 4.2% lower than ELEMENT-OF’s 64.3% which suggest the ELEMENT-OF corpus is more effective when the number of training examples is similar. When compared with PD-CROWD, the same amount of the gold annotated single-antecedent coreference examples (6k) achieved broadly the same score. Adding more training examples results in an increase in lenient scores. The strict scores follow a similar trend up until 2/3 of the examples are used (20k). Overall, the auxiliary corpus size is an important factor for the final results.

6 Other Approaches to Split-Antecedent Anaphora Resolution

Recently, Vala et al. (2016) introduced the first modern system to resolve split antecedent anaphora, although focusing only on plural pronouns *they* and *them*, and using a corpus of fiction they themselves annotated. Vala et al. (2016) proposed a learning-based system using handcrafted features, which achieved a score of 43.4% using the lenient evaluation they proposed and we adopted. The version of the task tackled in this paper is harder in three respects. First, our system resolves all split-antecedent references, without restriction. Secondly, our system was evaluated on the full ARRAU corpus (Uryupina et al., 2020), which contains text from multiple genres (news, dialogues, stories, medical and art history).

BEST	The sudden romance of British Aerospace and Thomson-CSF – traditionally bitter competitors for Middle East and Third World weapons contracts – is stirring controversy in Western Europe’s defense industry . Most threatened by closer <u>British Aerospace-Thomson</u> ties would be their respective national rivals
BASELINE	The sudden romance of British Aerospace and Thomson-CSF – traditionally bitter competitors for Middle East and Third World weapons contracts – is stirring controversy in Western Europe’s defense industry . Most threatened by closer <u>British Aerospace-Thomson</u> ties would be their respective national rivals
BEST	Workers dumped large burlap sacks of the imported material into a huge bin , poured in cotton and acetate fibers and mechanically mixed <u>the dry fibers</u> in a process used to make filters .
BASELINE	Workers dumped large burlap sacks of the imported material into a huge bin , poured in cotton and acetate fibers and mechanically mixed <u>the dry fibers</u> in a process used to make filters .
BEST	Time Warner Inc. is considering a legal challenge to Tele-Communications Inc. ’s plan to buy half of Showtime Networks Inc. , a move that could lead to all-out war between <u>the cable industry’s two most powerful players</u> .
BASELINE	Time Warner Inc. is considering a legal challenge to Tele-Communications Inc. ’s plan to buy half of Showtime Networks Inc. , a move that could lead to all-out war between <u>the cable industry’s two most powerful players</u> .
BEST	In California and New York , state officials have opposed Channel One . Mr. Whittle said private and parochial schools in <u>both states</u> will be canvassed to see if they are interested in ...
BASELINE	In California and New York , state officials have opposed Channel One . Mr. Whittle said private and parochial schools in <u>both states</u> will be canvassed to see if they are interested in ...

Table 5: A comparison of system prediction examples from our BEST and BASELINE system. The colours indicate the correctness of the predicted split-antecedents (**true positive** , **false negative** , **false positive**); the anchors are marked with underlines.

Thirdly, in addition to Vala *et al.*’s lenient evaluation that gives partial credit to split-antecedent anaphors not all of whose antecedents are identified, we also report strict scores that only give credit to the model when all the antecedents of an anaphor are correctly resolved.

More recently, Zhou and Choi (2018) introduced a corpus for entity linking and coreference in transcripts of the *Friends* sitcom. Plural mentions are annotated if they are linked to the main characters; as a result, the vast majority (95%) of the plurals in this corpus are pronouns. And since the corpus was primarily created for entity linking, its plural annotations are problematic for coreference: 58.8% of plural mentions are linked either to *General* entities that are not annotated in the text, or to characters that do not appear in the utterance before the plural anaphor. Also, only the results on the combination of singular and plural mentions are reported; the performance on plural mentions only is not.

7 Conclusions

We propose the first model for unrestricted split-antecedent anaphora resolution. Starting from the SoTA single-antecedent coreference resolution system, we substantially improve its performance on the task through a combination of exploiting auxiliary corpora for related tasks. Despite our baseline having a good performance of 58.6% (lenient F1), our best model achieved large gains of 11 percentage points. Further, evaluation using strict accuracy shows our best system can correctly resolve 43.6% of split-antecedent anaphors, which is 21 p.p. better than our baseline.

Acknowledgements

This research was supported in part by the DALI project, ERC Grant 695662.

References

- Anders Björkelund and Jonas Kuhn. 2014. Learning structured perceptrons for coreference resolution with latent antecedents and non-local features. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 47–57, Baltimore, Maryland, June. Association for Computational Linguistics.
- Kevin Clark and Christopher D. Manning. 2015. Entity-centric coreference resolution with model stacking. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1405–1415, Beijing, China, July. Association for Computational Linguistics.
- Kevin Clark and Christopher D. Manning. 2016. Improving coreference resolution by learning entity-level distributed representations. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 643–653, Berlin, Germany, August. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Urvashi Khandelwal, Christopher D. Manning, and Quoc V. Le. 2019. BAM! born-again multi-task networks for natural language understanding. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5931–5937, Florence, Italy, July. Association for Computational Linguistics.
- Ryan Cotterell and Kevin Duh. 2017. Low-resource named entity recognition with cross-lingual, character-level neural conditional random fields. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 91–96, Taipei, Taiwan, November. Asian Federation of Natural Language Processing.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. 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 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Carola Eschenbach, Christopher Habel, Michael Herweg, and Klaus Rehkämper. 1989. Remarks on plural anaphora. In *Proceedings of the fourth conference on European chapter of the Association for Computational Linguistics*, pages 161–167. Association for Computational Linguistics.
- Yufang Hou. 2020. Bridging anaphora resolution as question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1428–1438, Online, July. Association for Computational Linguistics.
- Mandar Joshi, Omer Levy, Luke Zettlemoyer, and Daniel Weld. 2019. BERT for coreference resolution: Baselines and analysis. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5803–5808, Hong Kong, China, November. Association for Computational Linguistics.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Hans Kamp and Uwe Reyle. 1993. *From Discourse to Logic*. D. Reidel, Dordrecht.
- Ben Kantor and Amir Globerson. 2019. Coreference resolution with entity equalization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 673–677, Florence, Italy, July. Association for Computational Linguistics.
- Barbara Kaup, Stephanie Kelter, and Christopher Habel. 2002. Representing referents of plural expressions and resolving plural anaphors. *Language and Cognitive Processes*, 17(4):405–450.
- Vid Kocijan, Oana-Maria Camburu, Ana-Maria Cretu, Yordan Yordanov, Phil Blunsom, and Thomas Lukasiewicz. 2019. WikiCREM: A large unsupervised corpus for coreference resolution. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4303–4312, Hong Kong, China, November. Association for Computational Linguistics.
- Ekaterina Lapshinova-Koltunski, Christian Hardmeier, and Pauline Krielke. 2018. ParCorFull: a parallel corpus annotated with full coreference. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, pages 423–428, Miyazaki, Japan, May. European Language Resources Association (ELRA).

- Heeyoung. Lee, Angel. Chang, Yves. Peirsman, Nathaneal. Chambers, Mihai. Surdeanu, and Dan. Jurafsky. 2013. Deterministic coreference resolution based on entity-centric, precision-ranked rules. *Computational Linguistics*, 39(4):885–916.
- 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*, pages 188–197, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. Higher-order coreference resolution with coarse-to-fine inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 687–692, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Gregory L. Murphy. 1984. Establishing and accessing referents in discourse. *Memory & Cognition*, 12(5):489–497.
- Nikole D. Patson. 2014. The processing of plural expressions. *Language and Linguistics Compass*, 8(8):319–329.
- Silviu Paun, Jon Chamberlain, Udo Kruschwitz, Juntao Yu, and Massimo Poesio. 2018. A probabilistic annotation model for crowdsourcing coreference. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1926–1937, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Viktor Pekar, Juntao Yu, Mohab El-karef, and Bernd Bohnet. 2014. Exploring options for fast domain adaptation of dependency parsers. In *Proceedings of the First Joint Workshop on Statistical Parsing of Morphologically Rich Languages and Syntactic Analysis of Non-Canonical Languages*, pages 54–65, Dublin, Ireland, August. Dublin City University.
- Haoruo Peng, Daniel Khashabi, and Dan Roth. 2015. Solving hard coreference problems. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 809–819, Denver, Colorado, May–June. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. 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 2227–2237, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Massimo Poesio and Ron Artstein. 2008. Anaphoric annotation in the ARRAU corpus. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC’08)*, Marrakech, Morocco, May. European Language Resources Association (ELRA).
- Massimo Poesio, Roland Stuckardt, and Yannick Versley. 2016. *Anaphora Resolution: Algorithms, Resources and Applications*. Springer, Berlin.
- Massimo Poesio, Jon Chamberlain, Silviu Paun, Juntao Yu, Alexandra Uma, and Udo Kruschwitz. 2019. A crowdsourced corpus of multiple judgments and disagreement on anaphoric interpretation. 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 1778–1789, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes. In *Proceedings of the Sixteenth Conference on Computational Natural Language Learning (CoNLL 2012)*, Jeju, Korea.
- Altaf Rahman and Vincent Ng. 2012. Resolving complex cases of definite pronouns: The Winograd schema challenge. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 777–789, Jeju Island, Korea, July. Association for Computational Linguistics.
- Anthony J. Sanford and F. Lockhart. 1990. Description types and method of conjoining as factors influencing plural anaphora. *Journal of Semantics*, 7:365–378.

- Wee M. Soon, Daniel C. Y. Lim, and Hwee T. Ng. 2001. A machine learning approach to coreference resolution of noun phrases. *Computational Linguistics*, 27(4).
- Olga Uryupina, Ron Artstein, Antonella Bristot, Federica Cavicchio, Francesca Delogu, Kepa J. Rodriguez, and Massimo Poesio. 2020. Annotating a broad range of anaphoric phenomena, in a variety of genres: the ARRAU corpus. *Journal of Natural Language Engineering*.
- Hardik Vala, Andrew Piper, and Derek Ruths. 2016. The more antecedents, the merrier: Resolving multi-antecedent anaphors. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2287–2296, Berlin, Germany, August. Association for Computational Linguistics.
- Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. 2018. Mind the GAP: A balanced corpus of gendered ambiguous pronouns. *Transactions of the Association for Computational Linguistics*, 6:605–617.
- Sam Wiseman, Alexander M. Rush, Stuart Shieber, and Jason Weston. 2015. Learning anaphoricity and antecedent ranking features for coreference resolution. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1416–1426, Beijing, China, July. Association for Computational Linguistics.
- Sam Wiseman, Alexander M. Rush, and Stuart M. Shieber. 2016. Learning global features for coreference resolution. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 994–1004, San Diego, California, June. Association for Computational Linguistics.
- Zhilin Yang, Ruslan Salakhutdinov, and William W Cohen. 2017. Transfer learning for sequence tagging with hierarchical recurrent networks. *Proceedings of the ICLR*.
- Juntao Yu and Bernd Bohnet. 2015. Exploring confidence-based self-training for multilingual dependency parsing in an under-resourced language scenario. In *Proceedings of the Third International Conference on Dependency Linguistics (Depling 2015)*, pages 350–358, Uppsala, Sweden, August. Uppsala University, Uppsala, Sweden.
- Juntao Yu, Mohab Elkaref, and Bernd Bohnet. 2015. Domain adaptation for dependency parsing via self-training. In *Proceedings of the 14th International Conference on Parsing Technologies*, pages 1–10, Bilbao, Spain, July. Association for Computational Linguistics.
- Ethan Zhou and Jinho D. Choi. 2018. They exist! introducing plural mentions to coreference resolution and entity linking. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 24–34, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Joey Tianyi Zhou, Hao Zhang, Di Jin, Hongyuan Zhu, Meng Fang, Rick Siow Mong Goh, and Kenneth Kwok. 2019. Dual adversarial neural transfer for low-resource named entity recognition. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3461–3471, Florence, Italy, July. Association for Computational Linguistics.