

Ensemble Transfer Learning for Multilingual Coreference Resolution

Tuan Lai, Heng Ji

Department of Computer Science
University of Illinois Urbana-Champaign
{tuanml2, hengji}@illinois.edu

Abstract

Entity coreference resolution is an important research problem with many applications, including information extraction and question answering. Coreference resolution for English has been studied extensively. However, there is relatively little work for other languages. A problem that frequently occurs when working with a non-English language is the scarcity of annotated training data. To overcome this challenge, we design a simple but effective ensemble-based framework that combines various transfer learning (TL) techniques. We first train several models using different TL methods. Then, during inference, we compute the unweighted average scores of the models' predictions to extract the final set of predicted clusters. Furthermore, we also propose a low-cost TL method that bootstraps coreference resolution models by utilizing Wikipedia anchor texts. Leveraging the idea that the coreferential links naturally exist between anchor texts pointing to the same article, our method builds a sizeable distantly-supervised dataset for the target language that consists of tens of thousands of documents. We can pre-train a model on the pseudo-labeled dataset before finetuning it on the final target dataset. Experimental results on two benchmark datasets, OntoNotes and SemEval, confirm the effectiveness of our methods. Our best ensembles consistently outperform the baseline approach of simple training by up to 7.68% in the F1 score. These ensembles also achieve new state-of-the-art results for three languages: Arabic, Dutch, and Spanish¹.

1 Introduction

Within-document entity coreference resolution is the process of clustering entity mentions in a document that refer to the same entities (Ji et al., 2005; Luo and Zitouni, 2005; Ng, 2010, 2017). It is an important research problem, with applications in various downstream tasks such as entity linking

(Ling et al., 2015; Kundu et al., 2018), question answering (Dhingra et al., 2018), and dialog systems (Gao et al., 2019). Researchers have recently proposed many neural methods for coreference resolution, ranging from span-based end-to-end models (Lee et al., 2017, 2018) to formulating the task as a question answering problem (Wu et al., 2020b). Given enough annotated training data, deep neural networks can learn to extract useful features automatically. As a result, on English benchmarks with abundant labeled training documents, the mentioned neural methods consistently outperform previous handcrafted feature-based techniques (Raghuathan et al., 2010; Lee et al., 2013), achieving new state-of-the-art (SOTA) results.

Compared to the amount of research on English coreference resolution, there is relatively little work for other languages. A problem that frequently occurs when working with a non-English language is the scarcity of annotated training data. For example, the benchmark OntoNotes dataset contains about eight times more documents in English than in Arabic (Pradhan et al., 2012). Some recent studies aim to overcome this challenge by applying standard cross-lingual transfer learning (TL) methods such as continued training or joint training (Kundu et al., 2018; Pražák et al., 2021). In continued training, a model pretrained on a source dataset is further finetuned on a (typically smaller) target dataset (Xia and Van Durme, 2021). In joint training, a model is trained on the concatenation of the source and target datasets (Min, 2021). The mentioned studies only use one transfer method at a time, and they do not explore how to combine multiple TL techniques effectively. This can be sub-optimal since different learning methods can be complementary (Liu et al., 2019; Li et al., 2021). For example, our experimental results to be discussed later show that continued training and joint training are highly complementary. Furthermore, a disadvantage of using a cross-lingual transfer method is the require-

¹Data and code will be made available upon publication.

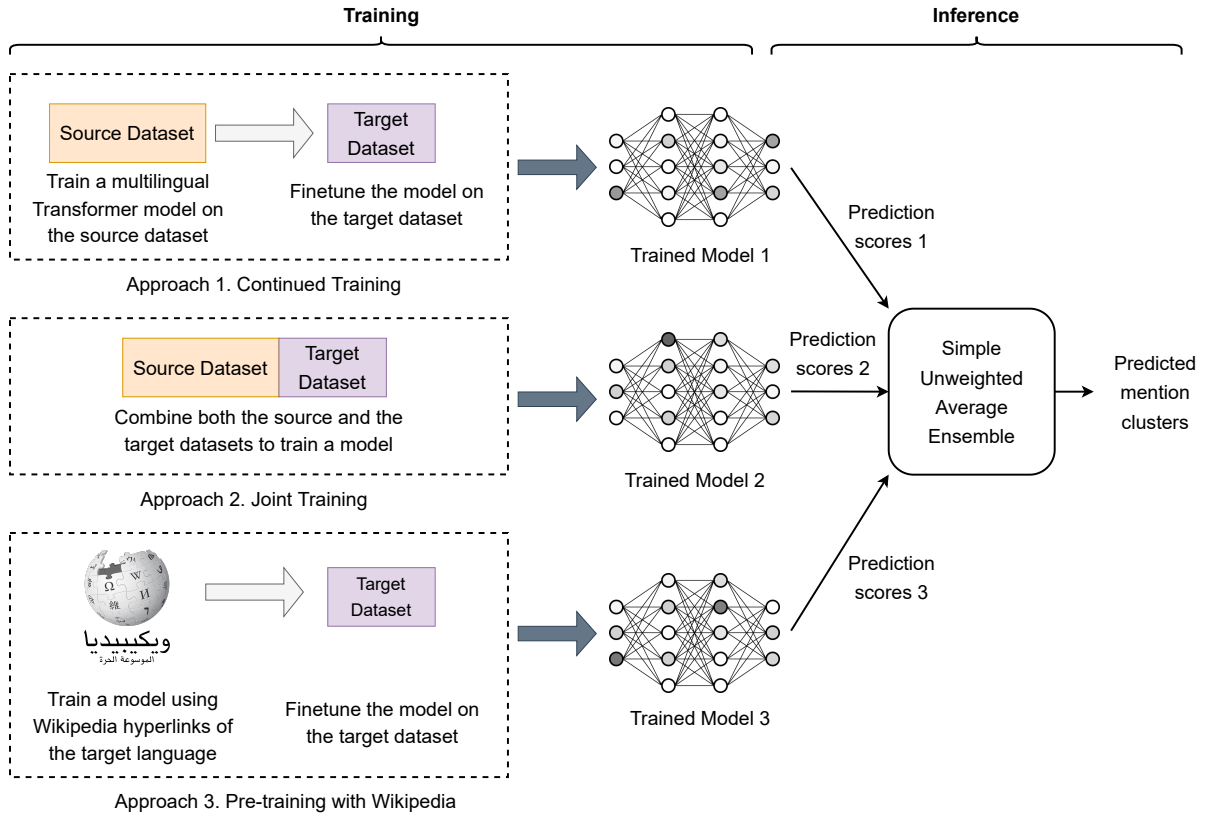


Figure 1: An overview of our framework. We first train several coreference resolution models using different TL approaches. During inference, we use a simple unweighted averaging method to combine the models’ predictions.

ment of a labeled coreference resolution dataset in some source language (usually English).

In this work, we propose an effective ensemble-based framework for combining various TL techniques. We first train several coreference models using different TL methods. During inference, we compute the unweighted average scores of the models’ predictions to extract the final set of mention clusters. We also propose a low-cost TL method that bootstraps coreference models without using a labeled dataset in some source language. The basic idea is that the coreference relation often holds between anchor texts pointing to the same Wikipedia article. Based on this observation, our TL method builds a sizable distantly-supervised dataset for the target language from Wikipedia. We can then pre-train a model on the pseudo-labeled dataset before finetuning it on the final target dataset. Experimental results on two datasets, OntoNotes and SemEval (Recasens et al., 2010), confirm the effectiveness of our proposed methods. Our best ensembles outperform the baseline approach of simple training by up to 7.68% absolute gain in the F1 score. These ensembles also achieve new SOTA results for three languages: Arabic, Dutch, and Spanish.

In summary, our main contributions include:

- We introduce an ensemble-based framework that combines various TL methods effectively.
- We design a new TL method that leverages Wikipedia to bootstrap coreference models.
- Extensive experimental results show that our proposed methods are highly effective and provide useful insights into entity coreference resolution for non-English languages.

2 Methods

Figure 1 shows an overview of our framework. During the training stage, we train several coreference resolution models using various TL approaches. For simplicity, we use the same span-based architecture (Section 2.1) for every model to be trained. However, starting from the same architecture, using different learning methods typically results in models with different parameters. In this work, our framework uses two types of TL methods: (a) cross-lingual TL approaches (Section 2.2) and (b) our newly proposed Wikipedia-based approach (Section 2.3). The cross-lingual TL methods require a labeled coreference resolution dataset in some

source language, but our Wikipedia-based method does not have that limitation. Our framework is general as it can work with other learning methods (e.g., self-distillation). During inference, we use a simple unweighted averaging method to combine the trained models’ predictions (Section 2.4).

2.1 Span-based End-to-End Coreference Resolution

In this work, the architecture of every model is based on the popular span-based *e2e-coref* model (Lee et al., 2017). Given an input document consisting of n tokens, our model first forms a contextualized representation for each input token using a multilingual Transformer encoder such as XLM-R (Conneau et al., 2020). Let $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ be the output of the encoder. For each candidate span i , we define its representation \mathbf{g}_i as:

$$\mathbf{g}_i = [\mathbf{x}_{\text{START}(i)}, \mathbf{x}_{\text{END}(i)}, \hat{\mathbf{x}}_i, \phi(s_i)] \quad (1)$$

where $\text{START}(i)$ and $\text{END}(i)$ denote the start and end indices of span i respectively. $\hat{\mathbf{x}}_i$ is an attention-weighted sum of the token representations in the span (Lee et al., 2017). $\phi(s_i)$ is a feature vector encoding the size of the span.

To maintain tractability, we only consider spans with up to L tokens. The value of L is selected empirically and set to be 30. All the span representations are fed into a mention scorer $s_m(\cdot)$:

$$s_m(i) = \text{FFNN}_m(\mathbf{g}_i) \quad (2)$$

where FFNN_m is a feedforward neural network with ReLU activations. Intuitively, $s_m(i)$ indicates whether span i is indeed an entity mention.

After scoring the spans using FFNN_m , we only keep spans with high mention scores². We denote the set of the unpruned spans as S . Then, for each remaining span $i \in S$, the model predicts a distribution $\hat{P}(j)$ over its antecedents³ $j \in Y(i)$:

$$\hat{P}(j) = \frac{\exp(s(i, j))}{\sum_{k \in Y(i)} \exp(s(i, k))} \quad (3)$$

$$s(i, j) = \text{FFNN}_s([\mathbf{g}_i, \mathbf{g}_j, \mathbf{g}_i \circ \mathbf{g}_j, \phi(i, j)])$$

where $Y(i) = \{\epsilon, 1, \dots, i - 1\}$ is a set consisting of a dummy antecedent ϵ and all spans that precede i . The dummy antecedent ϵ represents two possible

cases: (1) the span i is not an entity mention, or (2) the span i is an entity mention, but it is not coreferential with any remaining preceding span. FFNN_s is a feedforward network, and \circ is element-wise multiplication. $\phi(i, j)$ encodes the distance between the two spans i and j . Finally, note that $s(i, \epsilon)$ is fixed to be 0.

Given a labeled document D and a model with parameters θ , we define the mention detection loss:

$$\mathcal{L}_{\text{detect}}(\theta, D) = -\frac{1}{|S|} \sum_{i \in S} \mathcal{L}_{\text{detect}}(\theta, i)$$

$$\mathcal{L}_{\text{detect}}(\theta, i) = y_i \log \hat{y}_i + (1 - y_i) \log (1 - \hat{y}_i)$$

where $\hat{y}_i = \text{sigmoid}(s_m(i))$, and $y_i = 1$ if and only if span i is in one of the gold-standard mention clusters. In addition, we also want to maximize the marginal log-likelihood of all correct antecedents implied by the gold-standard clustering:

$$\mathcal{L}_{\text{cluster}}(\theta, D) = -\log \prod_{i \in S} \sum_{\hat{y} \in Y(i) \cap \text{GOLD}(i)} \hat{P}(\hat{y})$$

where $\text{GOLD}(i)$ are gold antecedents for span i . $\hat{P}(\hat{y})$ is calculated using Equation 3. Our final loss combines mention detection and clustering:

$$\mathcal{L}(\theta, D) = \mathcal{L}_{\text{detect}}(\theta, D) + \mathcal{L}_{\text{cluster}}(\theta, D) \quad (4)$$

2.2 Cross-Lingual Transfer Learning

Inspired by previous studies (Xia and Van Durme, 2021; Min, 2021; Pražák et al., 2021), we investigate two different cross-lingual transfer learning methods: *continued training* and *joint training*. Both methods assume the existence of a labeled dataset in some source language. In this work, we use the English OntoNotes dataset (Pradhan et al., 2012) as the source dataset, as it contains nearly 3,500 annotated documents (Table 1).

Continued Training. We first train a coreference resolution model on the source dataset until convergence. After that, we further finetune the pre-trained model on a target dataset. More formally, let $M(f, \theta_0)$ denote an optimization procedure for f with initial guess θ_0 . This optimization procedure can, for example, be the application of some stochastic gradient descent algorithm. Also, let \mathbb{S} be the set of all training documents in the source dataset, and let \mathbb{T} denote the set of all training documents in the target dataset. Then, the first stage of continued training can be described as:

$$\hat{\theta}_1 = M\left(\sum_{D \in \mathbb{S}} \mathcal{L}(\theta, D), \hat{\theta}_0\right) \quad (5)$$

²We describe the exact filtering criteria in Section 3.1.

³All spans are ordered based on their start indices. Spans with the same start index are ordered by their end indices.

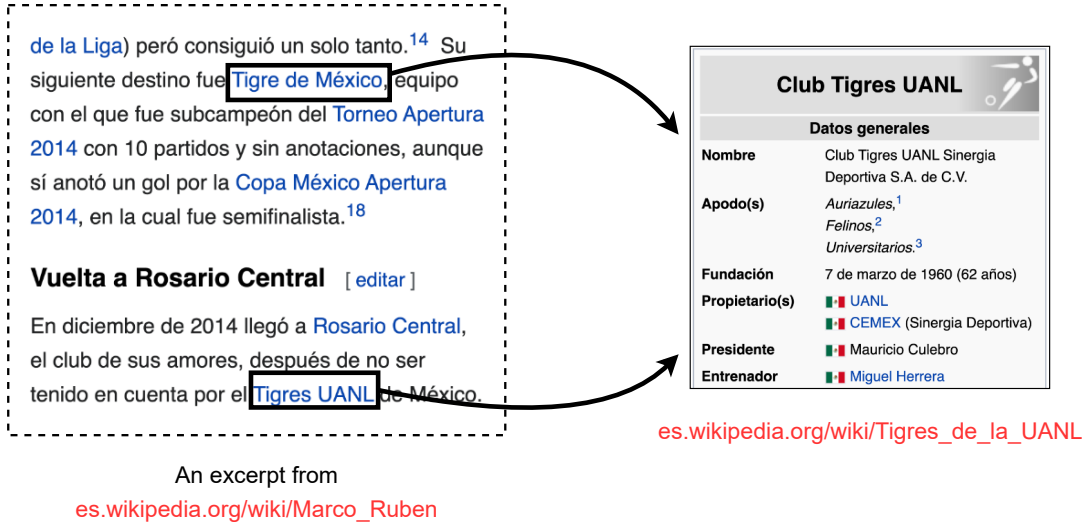


Figure 2: Since the hyperlinks of *Tigre de México* and *Tigres UANL* point to the same Wikipedia page, a person who does not know Spanish can still guess that the two mentions are likely to be coreferential. In fact, the two mentions both refer to Tigres UANL, a Mexican professional football club.

where $\hat{\theta}_0$ is randomly initialized. Then, the second stage can be described using Equation 6:

$$\hat{\theta}_2 = M\left(\sum_{D \in \mathbb{T}} \mathcal{L}(\theta, D), \hat{\theta}_1\right) \quad (6)$$

Here, $\hat{\theta}_2$ is the parameter set of the final model.

Joint Training. We combine both the source and target datasets to train a model. More specifically, using the same notations as above, we can describe joint training by the following equation:

$$\hat{\theta}_1 = M\left(\sum_{D \in \mathbb{T} \cup \mathbb{S}} \mathcal{L}(\theta, D), \hat{\theta}_0\right) \quad (7)$$

where $\hat{\theta}_0$ is randomly initialized, and $\hat{\theta}_1$ is the parameter set of the final model.

2.3 Bootstrapping using Wikipedia Hyperlinks

While cross-lingual methods such as continued training and joint training are conceptually simple and typically effective (Huang et al., 2020), they require the existence of a labeled dataset in some source language. To overcome this limitation, we propose an inexpensive TL method that bootstraps coreference models by utilizing Wikipedia anchor texts. The basic idea is that two anchor texts pointing to the same Wikipedia page are likely coreferential (See Figure 2 for an example). Our method builds a large distantly-supervised dataset \mathbb{W} for

the target language by leveraging this observation:

$$\mathbb{W} = \{D_1, D_2, \dots, D_m\} \quad (8)$$

where D_i is a text document constructed from some Wikipedia page written in the target language. The number of mentions in D_i is the same as the number of anchor texts in the text portion of the original Wikipedia article. We consider two mentions in D_i to be coreferential if and only if their corresponding anchor texts point to the same article.

After constructing \mathbb{W} , we follow a two-step process similar to the continued training approach. We first train a coreference resolution model on \mathbb{W} until convergence. Then, we finetune the pre-trained model on the final target dataset.

Compared to a manually-labeled dataset, \mathbb{W} has several disadvantages. Not all entity mentions are exhaustively marked in Wikipedia documents. For example, in Spanish Wikipedia, pronouns are typically not annotated. Nevertheless, since Wikipedia is one of the largest multilingual repositories for information, \mathbb{W} is generally large (see Table 1 for some statistics), and it contains documents on various topics. As such, \mathbb{W} can still provide some useful distant supervision signals, and so it can serve as a source dataset in the TL process.

2.4 Ensemble-Based Coreference Resolution

During the training stage, we train three different coreference resolution models using the TL approaches described above. At test time, we use a simple unweighted averaging method to combine

the models’ predictions. More specifically, for a candidate span i with no more than L tokens, we compute its mention score as follows:

$$s_{m,\text{ensemble}}(i) = \frac{(s_{m,1}(i) + s_{m,2}(i) + s_{m,3}(i))}{3} \quad (9)$$

where $s_{m,1}(i)$, $s_{m,2}(i)$, and $s_{m,3}(i)$ are the mention scores produced by the three trained models separately (refer to Equation 2). Intuitively, these scores indicate whether span i is an entity mention.

Similar to the process described in Section 2.1, after scoring every span whose length is no more than L using Equation 9, we only keep spans with high mention scores⁴. Then, for each remaining span i , we predict a distribution over its antecedents $j \in Y(i)$ as follows:

$$\hat{P}_{\text{ensemble}}(j) = \frac{\exp(s_{\text{ensemble}}(i, j))}{\sum_{k \in Y(i)} \exp(s_{\text{ensemble}}(i, k))}$$

$$s_{\text{ensemble}}(i, j) = \frac{s_1(i, j) + s_2(i, j) + s_3(i, j)}{3} \quad (10)$$

where $s_1(i, j)$, $s_2(i, j)$, and $s_3(i, j)$ are the pairwise scores produced by the trained models separately (Eq. 3). We fix $s_{\text{ensemble}}(i, \epsilon)$ to be 0.

After computing the antecedent distribution for each remaining span, we can extract the final set of mention clusters. Note that while we consider only three individual TL methods in this work, Equation 9 and Equation 10 can easily be extended for the case when we use more TL methods.

3 Experiments

3.1 Data and Experiments Setup

Evaluation metrics Following prior work (Pradhan et al., 2012), we evaluate coreference using the average F_1 between B^3 (Bagga and Baldwin, 1998), MUC (Vilain et al., 1995), and CEAF _{ϕ_4} (Luo, 2005). We refer to this metric as AVG.

Datasets Table 1 shows the basic statistics of all the datasets we used in this work. When using a cross-lingual TL method, we use the English portion of OntoNotes (Pradhan et al., 2012) as the source dataset. We explore three target datasets: OntoNotes Arabic (Pradhan et al., 2012), SemEval Dutch (Recasens et al., 2010), and SemEval Spanish (Recasens et al., 2010). These datasets contain data in three different languages.

⁴We describe the exact filtering criteria in Section 3.1.

Dataset	Training	Dev	Test
<i>Source Datasets</i>			
OntoNotes English	2,802	343	348
Wikipedia-based Arabic Dataset	64,850	250	250
Wikipedia-based Dutch Dataset	46,715	250	250
Wikipedia-based Spanish Dataset	104,520	250	250
<i>Target Datasets</i>			
OntoNotes Arabic	359	44	44
SemEval Dutch	145	23	72
SemEval Spanish	875	140	168

Table 1: Number of documents for each of the datasets.

OntoNotes does not annotate singleton mentions (i.e., noun phrases not involved in any coreference chain). It only has annotations for non-singleton mentions. SemEval has annotations for singletons.

Wikipedia-based Dataset Construction To construct a distantly-supervised dataset, we first download a complete Wikipedia dump in the target language. We then extract clean text and hyperlinks from the dump using WikiExtractor⁵. For each pre-processed article, we cluster its anchor texts based on the destinations of their hyperlinks. We also filter out articles with too few coreference links (e.g., articles that only have singleton mentions).

General Hyperparameters We use two different learning rates, one for the lower pretrained Transformer encoder and one for the upper layers. For every setting, the lower learning rate is 1e-5, the upper learning rate is 1e-4, and the span length limit L is 30. The number of training/pre-training epochs is set to be 25 in most cases. When pre-training a model on a Wikipedia-based dataset, the number of epochs is 5. When fine-tuning a model already pre-trained on Dutch Wikipedia or Spanish Wikipedia, the number of epochs is 50. During each training/pre-training process, we pick the checkpoint which achieves the best AVG score on the appropriate dev set as the final checkpoint.

Transformer Encoders When the target dataset is OntoNotes Arabic, we use GigaBERT (Lan et al., 2020) as the Transformer encoder. GigaBERT is an English-Arabic bilingual language model pre-trained from the English and Arabic Gigaword corpora. When the target dataset is SemEval Dutch or SemEval Spanish, we use the multilingual XLM-RoBERTa (XLM-R) Transformer model (Conneau et al., 2020). More specifically, we use the base version of XLM-R (i.e., xlm-roberta-base).

⁵<https://tinyurl.com/wikiextractor>

	Arabic	Dutch	Spanish
<i>Baselines</i>			
◆ Previous SOTA (Table 3)	64.55	55.40	51.30
◆ Baseline Approach (trained using the target dataset)	63.70	52.81	72.18
<i>Individual Transfer/Pretraining Methods (Sections 2.2 and 2.3)</i>			
■ Continued Training	64.96	58.90	74.05
■ Joint Training	65.50	58.76	73.53
■ Wikipedia Pre-Training	63.78	53.15	73.35
<i>Ensembles (Section 2.4)</i>			
◆ Three models, each trained using the baseline approach	64.70	54.44	73.35
◆ Baseline Approach \oplus Wikipedia Pre-Training	65.75	55.25	74.19
◆ Joint Training \oplus Wikipedia Pre-Training	66.63	58.18	74.82
◆ Continued Training \oplus Wikipedia Pre-Training	66.24	57.88	75.43
◆ Continued Training \oplus Joint Training	65.79	60.49	74.93
◆ Continued Training \oplus Joint Training \oplus Wikipedia Pre-Training	66.72	59.66	75.62
<i>Oracle-Guided Ensembles (Section 3.3.1)</i>			
◇ Continued Training \oplus Joint Training \oplus Wikipedia Pre-Training	77.53	75.19	83.12

Table 2: Overall F_1 (in %) on OntoNotes Arabic, SemEval Dutch, and SemEval Spanish.

Dataset	Prior Work	Approach	Prev. Score	Our Best
OntoNotes Arabic	(Min, 2021)	GigaBERT + C2F + Joint Training	64.55	66.72
SemEval Dutch	(Xia and Van Durme, 2021)	XLM-R + ICoref + Continued Training	55.40	60.49
SemEval Spanish	(Xia and Van Durme, 2021)	XLM-R + ICoref + Continued Training	51.30	75.62

Table 3: Test F_1 (in %) on the target datasets and the previous SOTA on each dataset (to the best of our knowledge).

Span Pruning As described in Section 2.1, after computing a mention score for each span whose length is not more than L , we only keep spans with high scores. More specifically, when working with a dataset from OntoNotes (e.g., OntoNotes Arabic), we only keep up to λn spans with the highest mention scores (Lee et al., 2017). The value of λ is selected empirically and set to be 0.18. When working with any other dataset, we keep every span that has a positive mention score.

3.2 Overall Results

Table 2 shows the overall performance of different approaches. Our baseline approach is to simply train a model with the architecture described in Section 2.1 using only the target dataset of interest. Overall, the performance of a model trained using the baseline approach is positively correlated with the size of the corresponding target dataset, which is expected. A surprising finding is that our baseline approach already outperforms the previous SOTA method for SemEval Spanish (Xia and

Van Durme, 2021) by 20.98% in the F1 score. We speculate that the previous SOTA model for SemEval Spanish is severely undertrained.

Table 2 also shows the results of using different TL methods individually. Each of the TL methods can help improve the coreference resolution performance. While continued training seems to be the most effective approach, it requires the existence of a source dataset (OntoNotes English in this case). On the other hand, our newly proposed Wikipedia-based method can help improve the performance without relying on any labeled source dataset.

Finally, Table 2 also shows the results of using different combinations of learning approaches. Our simple unweighted averaging method is effective across almost all model combinations. In particular, by combining all of the three TL methods discussed previously, we can outperform the previous SOTA methods by large margins. In addition, even without using any labeled source dataset, the combination [Baseline Approach \oplus Wikipedia Pre-Training] can still outperform the previous SOTA

Approaches	AVG
Baseline Approach	40.10
Wikipedia Pre-Training	42.67
Baseline Approach \oplus Wikipedia Pre-Training	45.28

Table 4: Test F-score (in %) of various approaches on OntoNotes Arabic when we restrict the size of the gold Arabic training dataset to only 10 documents.

methods for Arabic and Spanish. This further confirms the usefulness of our Wikipedia-based TL method. Lastly, combining three models trained using the same baseline approach leads to smaller gains than combining the three TL methods. This is expected as ensemble methods typically work best when the individual learners are diverse (Krogh and Vedelsby, 1994; Melville and Mooney, 2003).

3.3 Analysis

3.3.1 How optimal is our simple unweighted averaging method?

Our averaging approach is equivalent to linear interpolation with equal weights. To analyze the optimality of our method, we compare it to the “best possible” interpolation method.

More specifically, we assume that there is an oracle that can tell us which model in an ensemble gives the most accurate prediction for a particular latent variable. Then, for example, suppose we want to score a span i using an ensemble of three models. If i is an entity mention, the oracle will tell us that the model that returns the highest mention score for i is the most accurate. Thus, we can set the score for i to be $\max(s_{m,1}(i), s_{m,2}(i), s_{m,3}(i))$. Following the same logic, if i is not an entity mention, we will set its score to be $\min(s_{m,1}(i), s_{m,2}(i), s_{m,3}(i))$. The same idea can be applied to compute the linking score $s_{\text{ensemble}}(i, j)$ between i and j .

In Table 2, we see a considerable gap between the performance of our simple averaging method and the oracle-guided interpolation method. Therefore, a promising future direction is to experiment with a more context-dependent ensemble method. Nevertheless, our averaging method is simple, and it does not require any further parameter tuning to combine a set of existing models. Finally, the performance of each oracle-guided ensemble is far from perfect, implying that improving the underlying architecture of each model can also be a worthwhile effort.

... el director del centro, Anna Mas, asegurar que el acto pretender “rechazar el agresión y concienciar a el alumno del incremento de este ataque”. el director recordar otro dos agresión “por llevar hierro dental o el pelo largo”. en mucho ocasión se producir asalto a niño, y el alumno, añadir Mas, “ver como algo normal que les parir por el calle y les quitar el poco dinero que llevar encima” ...

... Het vakblad Hormones and Behavior beschrijven hoe het voldoende zijn dat mannelijk muis een vleug vrouw roken om hen weinig bang te maken van kat en wezel ...

Table 5: Examples of mention clusters that were correctly predicted by our ensembles. Blue spans represent coreferential mentions. The first example is in Spanish. The second example is in Dutch.

3.3.2 How effective is our framework in extremely low-resource settings?

We conduct experiments on OntoNotes Arabic where we assume that the training dataset for Arabic only has 10 documents and that we do not have any source dataset (Table 4). In this setting, our ensemble substantially outperforms the baseline approach by up to 5.18% in the F1 score.

3.3.3 Qualitative Analysis

We provide some qualitative analyses to demonstrate the strengths of our ensembles in Table 5.

In the first example, the three highlighted mentions refer to Anna Mas, the director of a center. Our model trained using joint training merged this cluster with a different cluster that refers to a different entity (not shown in the example because of space constraints). In contrast, our models trained using other TL methods did not make that error. As a result, our best ensemble for Spanish predicted the correct cluster for Anna Mas.

The second example is in Dutch. Here, *mannelijk muis* can be translated as *male mouses*, while *hen* can be translated as *them*. Our model trained using continued training failed to extract the mention *mannelijk muis*. Nevertheless, in the end, our ensemble for Dutch was able to extract the mention and correctly link it to the pronoun *hen*.

4 Related Work

4.1 Entity Coreference Resolution

Recently, neural models for entity coreference resolution have shown superior performance over approaches using hand-crafted features. Lee et al. (2017) proposed the first end-to-end neural coreference resolution model named *e2e-coref*. The model uses a bi-directional LSTM and a head-finding attention mechanism to learn mention representations and calculate mention and antecedent scores. Lee

et al. (2018) extended the *e2e-coref* model by introducing a coarse-to-fine pruning mechanism and a higher-order inference mechanism. The model uses ELMo representations (Peters et al., 2018) instead of traditional word embeddings. The model is typically referred to as the *c2f-coref* model.

Almost all recent studies on entity coreference resolution are influenced by the design of *c2f-coref*. Joshi et al. (2019) built the *c2f-coref* system on top of BERT representations (Devlin et al., 2019). Fei et al. (2019) transformed *c2f-coref* into a policy gradient model that can optimize coreference evaluation metrics directly. Xu and Choi (2020) studied in depth the higher-order inference (HOI) mechanism of *c2f-coref*. The authors concluded that given a high-performing encoder such as SpanBERT (Joshi et al., 2020), the impact of HOI is negative to marginal. Another line of work aims to simplify and/or reduce the computational complexity of *c2f-coref* (Xia et al., 2020; Kirstain et al., 2021; Lai et al., 2021; Dobrovolskii, 2021).

The studies mentioned above only trained and evaluated models using English datasets such as OntoNotes English (Pradhan et al., 2012) and the GAP dataset (Webster et al., 2018). On the other hand, there is significantly less work on coreference resolution for other languages. For example, while *e2e-coref* was introduced in 2017, the first neural coreference resolver for Arabic was only recently proposed in 2020 (Aloraini et al., 2020). For Dutch, many existing systems are still using rule-based (van Cranenburgh, 2019) or traditional learning-based approaches (Hendrickx et al., 2008; De Clercq et al., 2011). Recently, Poot and van Cranenburgh (2020) evaluated the performance of *c2f-coref* on Dutch datasets of two different domains: literary novels and news/Wikipedia text.

While our models’ architecture is based on *e2e-coref* (Section 2.1), we go beyond just applying the models to a non-English language in this work. We propose new TL approaches that can take advantage of existing source datasets and Wikipedia to improve the final performance.

4.2 Transfer Learning for Coreference Resolution

Compared to English datasets, the size of a coreference resolution dataset for a non-English language is typically smaller. Several recent studies aim to overcome this challenge by applying standard cross-lingual TL methods such as continued train-

ing or joint training (Kundu et al., 2018; Xia and Van Durme, 2021; Pražák et al., 2021; Min, 2021). These studies only use one transfer method at a time, and they do not explore how to combine multiple TL techniques effectively. Our experimental results (Section 3.2) show that combining various TL techniques can substantially improve the final coreference resolution performance.

A closely related work by Yang et al. (2012) proposed an adaptive ensemble method to adapt coreference resolution across domains. Their study did not explicitly focus on improving coreference resolution for non-English languages. In addition, they experimented with the settings where gold standard mentions are assumed to be provided. We do not make that assumption. Each of our models does both mention extraction and linking.

4.3 Leveraging Wikipedia for Coreference Resolution

There have been studies on leveraging Wikipedia for coreference resolution. Eirew et al. (2021) recently created a large-scale cross-document event coreference dataset from English Wikipedia. For cross-document entity coreference, Singh et al. (2012) created Wikilinks by finding hyperlinks to English Wikipedia from a web crawl and using anchor text as mentions. Different from these studies, we focus on within-document entity coreference resolution. In addition, we explore coreference resolution for languages beyond English in this work.

Many previous studies leveraged Wikipedia for related tasks such as name tagging (Alotaibi and Lee, 2012; Nothman et al., 2013; Althobaiti et al., 2014) and entity linking (Pan et al., 2017; Wu et al., 2020a; Cao et al., 2021). We leave the extension of our methods to these tasks for future research.

5 Conclusions and Future Work

In this work, we propose an ensemble-based framework that combines various TL techniques. We also introduce a low-cost Wikipedia-based TL approach that does not require any labeled source dataset. Our approaches are highly effective, as our best ensembles achieve new SOTA results for three different languages. An interesting future direction is to explore the use of model compression techniques (Hinton et al., 2015; Han et al., 2016; Lai et al., 2020) to reduce the computational complexity of our ensembles.

6 Limitations

Multilingual language models such as XLM-R (Conneau et al., 2020) and GigaBERT (Lan et al., 2020) are typically pre-trained on large amounts of unlabeled text crawled from the Web. Since these models are optimized to capture the statistical properties of the training data, they tend to pick up on and amplify social stereotypes present in the data (Kurita et al., 2019). Since our coreference resolution models use such pre-trained language models, they may also exhibit social biases present on the Web. Identifying and mitigating social biases in neural models is an active area of research (Zhao et al., 2018; Sheng et al., 2021; Gupta et al., 2022). In the future, we plan to work on removing social biases from coreference resolution models.

Furthermore, while our proposed methods are highly effective, the performance of our best ensembles is still far from perfect. On OntoNotes Arabic, our best system only achieves an F1 score of 66.72%. Such performance may not be acceptable for some downstream tasks (e.g., information extraction from critical clinical notes).

Finally, even though Wikipedia is available in more than 300 languages, there are still very few Wikipedia pages for some very rare languages. Our proposed methods are likely to be less effective for such rare languages.

References

- Abdulrahman Aloraini, Juntao Yu, and Massimo Poesio. 2020. [Neural coreference resolution for Arabic](#). In *Proceedings of the Third Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 99–110, Barcelona, Spain (online). Association for Computational Linguistics.
- Fahd Alotaibi and Mark Lee. 2012. [Mapping Arabic Wikipedia into the named entities taxonomy](#). In *Proceedings of COLING 2012: Posters*, pages 43–52, Mumbai, India. The COLING 2012 Organizing Committee.
- Maha Althobaiti, Udo Kruschwitz, and Massimo Poesio. 2014. [Automatic creation of Arabic named entity annotated corpus using Wikipedia](#). In *Proceedings of the Student Research Workshop at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 106–115, Gothenburg, Sweden. Association for Computational Linguistics.
- Amit Bagga and Breck Baldwin. 1998. Algorithms for scoring coreference chains. In *The first international conference on language resources and evaluation workshop on linguistics coreference*, volume 1, pages 563–566. Citeseer.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. [Autoregressive entity retrieval](#). In *International Conference on Learning Representations*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *ACL*.
- Orphée De Clercq, Véronique Hoste, and Iris Hendrickx. 2011. [Cross-domain Dutch coreference resolution](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, pages 186–193, Hissar, Bulgaria. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Bhuwan Dhingra, Qiao Jin, Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. 2018. [Neural models for reasoning over multiple mentions using coreference](#). 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 42–48, New Orleans, Louisiana. Association for Computational Linguistics.
- Vladimir Dobrovolskii. 2021. [Word-level coreference resolution](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7670–7675, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alon Eirew, Arie Cattan, and Ido Dagan. 2021. [WEC: Deriving a large-scale cross-document event coreference dataset from Wikipedia](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2498–2510, Online. Association for Computational Linguistics.
- Hongliang Fei, Xu Li, Dingcheng Li, and Ping Li. 2019. [End-to-end deep reinforcement learning based coreference resolution](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 660–665, Florence, Italy. Association for Computational Linguistics.

- Yifan Gao, Piji Li, Irwin King, and Michael R. Lyu. 2019. [Interconnected question generation with coreference alignment and conversation flow modeling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4853–4862, Florence, Italy. Association for Computational Linguistics.
- Umang Gupta, Jwala Dhamala, Varun Kumar, Apurv Verma, Yada Pruksachatkun, Satyapriya Krishna, Rahul Gupta, Kai-Wei Chang, Greg Ver Steeg, and Aram Galstyan. 2022. Mitigating gender bias in distilled language models via counterfactual role reversal. In *ACL Finding*.
- Song Han, Huizi Mao, and William J. Dally. 2016. Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding. *arXiv: Computer Vision and Pattern Recognition*.
- Iris Hendrickx, Gosse Bouma, Frederik Coppens, Walter Daelemans, Veronique Hoste, Geert Kloosterman, Anne-Marie Mineur, Joeri Van Der Vloet, and Jean-Luc Verschelde. 2008. [A coreference corpus and resolution system for Dutch](#). In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the knowledge in a neural network. *ArXiv*, abs/1503.02531.
- Jocelyn Huang, Oleksii Kuchaiev, Patrick O’Neill, Vitaly Lavrukhin, Jason Li, Adriana Flores, Georg Kucsko, and Boris Ginsburg. 2020. Cross-language transfer learning, continuous learning, and domain adaptation for end-to-end automatic speech recognition. *arXiv preprint arXiv:2005.04290*.
- Heng Ji, David Westbrook, and Ralph Grishman. 2005. [Using semantic relations to refine coreference decisions](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 17–24, Vancouver, British Columbia, Canada. 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.
- 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. Association for Computational Linguistics.
- Yuval Kirstain, Ori Ram, and Omer Levy. 2021. [Coreference resolution without span representations](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 14–19, Online. Association for Computational Linguistics.
- Anders Krogh and Jesper Vedelsby. 1994. Neural network ensembles, cross validation, and active learning. *Advances in neural information processing systems*, 7.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*.
- Gourab Kundu, Avi Sil, Radu Florian, and Wael Hamza. 2018. [Neural cross-lingual coreference resolution and its application to entity linking](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 395–400, Melbourne, Australia. Association for Computational Linguistics.
- Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. [Measuring bias in contextualized word representations](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172, Florence, Italy. Association for Computational Linguistics.
- Tuan Manh Lai, Trung Bui, and Doo Soon Kim. 2021. End-to-end neural coreference resolution revisited: A simple yet effective baseline. *arXiv preprint arXiv:2107.01700*.
- Tuan Manh Lai, Quan Hung Tran, Trung Bui, and Daisuke Kihara. 2020. A simple but effective bert model for dialog state tracking on resource-limited systems. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8034–8038. IEEE.
- Wuwei Lan, Yang Chen, Wei Xu, and Alan Ritter. 2020. [An empirical study of pre-trained transformers for Arabic information extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4727–4734, Online. Association for Computational Linguistics.
- Heeyoung Lee, Angel Chang, Yves Peirsman, Nathanael 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. 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. Association for Computational Linguistics.
- Shiyang Li, Semih Yavuz, Wenhu Chen, and Xifeng Yan. 2021. [Task-adaptive pre-training and self-training are complementary for natural language understanding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1006–1015, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiao Ling, Sameer Singh, and Daniel S. Weld. 2015. [Design challenges for entity linking](#). *Transactions of the Association for Computational Linguistics*, 3:315–328.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019. Multi-task deep neural networks for natural language understanding. In *ACL*.
- Xiaoqiang Luo. 2005. On coreference resolution performance metrics. In *HLT*.
- Xiaoqiang Luo and Imed Zitouni. 2005. [Multi-lingual coreference resolution with syntactic features](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 660–667, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Prem Melville and Raymond J. Mooney. 2003. Constructing diverse classifier ensembles using artificial training examples. In *IJCAI*.
- Bonan Min. 2021. [Exploring pre-trained transformers and bilingual transfer learning for Arabic coreference resolution](#). In *Proceedings of the Fourth Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 94–99, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. 2020. [Pytorch metric learning](#).
- Vincent Ng. 2010. [Supervised noun phrase coreference research: The first fifteen years](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1396–1411, Uppsala, Sweden. Association for Computational Linguistics.
- Vincent Ng. 2017. Machine learning for entity coreference resolution: A retrospective look at two decades of research. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 4877–4884.
- Joel Nothman, Nicky Ringland, Will Radford, Tara Murphy, and James R. Curran. 2013. Learning multilingual named entity recognition from wikipedia. *Artif. Intell.*, 194:151–175.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958, Vancouver, Canada. Association for Computational Linguistics.
- Adam Paszke, S. Gross, Francisco Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, Alban Desmaison, Andreas Köpf, E. Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, B. Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*.
- Matthew E. 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. Association for Computational Linguistics.
- Corbèn Poot and Andreas van Cranenburgh. 2020. [A benchmark of rule-based and neural coreference resolution in Dutch novels and news](#). In *Proceedings of the Third Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 79–90, Barcelona, Spain (online). 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 *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea. Association for Computational Linguistics.
- Ondřej Pražák, Miloslav Konopík, and Jakub Sido. 2021. [Multilingual coreference resolution with harmonized annotations](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 1119–1123, Held Online. INCOMA Ltd.
- Karthik Raghunathan, Heeyoung Lee, Sudarshan Rangarajan, Nathanael Chambers, Mihai Surdeanu, Dan Jurafsky, and Christopher Manning. 2010. [A multi-pass sieve for coreference resolution](#). In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 492–501, Cambridge, MA. Association for Computational Linguistics.
- Marta Recasens, Lluís Màrquez, Emili Sapena, M. Antònia Martí, Mariona Taulé, Véronique Hoste, Massimo Poesio, and Yannick Versley. 2010. [SemEval-2010 task 1: Coreference resolution in multiple languages](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 1–8, Uppsala, Sweden. Association for Computational Linguistics.

- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2021. Societal biases in language generation: Progress and challenges. In *ACL*.
- Sameer Singh, Amarnag Subramanya, Fernando Pereira, and Andrew McCallum. 2012. Wikilinks: A large-scale cross-document coreference corpus labeled via links to wikipedia. *University of Massachusetts, Amherst, Tech. Rep. UM-CS-2012*, 15.
- Andreas van Cranenburgh. 2019. A dutch coreference resolution system with an evaluation on literary fiction. *Computational Linguistics in the Netherlands Journal*, 9:27–54.
- Marc B. Vilain, John D. Burger, John S. Aberdeen, Dennis Connolly, and Lynette Hirschman. 1995. A model-theoretic coreference scoring scheme. In *MUC*.
- 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.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2020a. **Scalable zero-shot entity linking with dense entity retrieval**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6397–6407, Online. Association for Computational Linguistics.
- Wei Wu, Fei Wang, Arianna Yuan, Fei Wu, and Jiwei Li. 2020b. **CorefQA: Coreference resolution as query-based span prediction**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6953–6963, Online. Association for Computational Linguistics.
- Patrick Xia, João Sedoc, and Benjamin Van Durme. 2020. **Incremental neural coreference resolution in constant memory**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8617–8624, Online. Association for Computational Linguistics.
- Patrick Xia and Benjamin Van Durme. 2021. **Moving on from OntoNotes: Coreference resolution model transfer**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5241–5256, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Liyang Xu and Jinho D. Choi. 2020. **Revealing the myth of higher-order inference in coreference resolution**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8527–8533, Online. Association for Computational Linguistics.
- Jian Bo Yang, Qi Mao, Qiao Liang Xiang, Ivor Wai-Hung Tsang, Kian Ming Adam Chai, and Hai Leong Chieu. 2012. **Domain adaptation for coreference resolution: An adaptive ensemble approach**. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 744–753, Jeju Island, Korea. Association for Computational Linguistics.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. **Gender bias in coreference resolution: Evaluation and debiasing methods**. 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 15–20, New Orleans, Louisiana. Association for Computational Linguistics.

A Reproducibility Information

In this section, we present the reproducibility information of our paper.

Implementation Dependencies Libraries Pytorch 1.11.0 (Paszke et al., 2019), Transformers 4.17.0 (Wolf et al., 2020), SentencePiece 0.1.96 (Kudo and Richardson, 2018), PyTorch Metric Learning (Musgrave et al., 2020).

Computing Infrastructure The experiments were conducted on a server with Intel(R) Xeon(R) Gold 5120 CPU @ 2.20GHz and NVIDIA Tesla V100 GPUs. GPU memory is 16G.

Number of Model Parameters When the target dataset is OntoNotes Arabic, we use GigaBERT (Lan et al., 2020) as the Transformer encoder. GigaBERT has about 125M parameters.

When the target dataset is SemEval Dutch or SemEval Spanish, we use the base version of XLM-R (i.e., xlm-roberta-base) (Conneau et al., 2020). xlm-roberta-base has about 278M parameters.

Hyperparameters The information about the hyperparameters is available in the main paper.

Expected Validation Performance We report the validation performance of the ensemble [*Continued Training* \oplus *Joint Training* \oplus *Wikipedia Pre-Training*].

The validation F1 score of the ensemble for Arabic coreference resolution is 66.60%. The total time needed for the evaluation is about 1 minute and 19 seconds.

The validation F1 score of the ensemble for Dutch coreference resolution is 57.81%. The total time needed for the evaluation is about 20 seconds.

The validation F1 score of the ensemble for Spanish coreference resolution is 75.73%. The total time needed for the evaluation is about 1 minute and 42 seconds.