



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2022

XCoref: Cross-document Coreference Resolution in the Wild

Zhukova, Anastasia ; Hamborg, Felix ; Donnay, Karsten ; Gipp, Bela

DOI: https://doi.org/10.1007/978-3-030-96957-8_25

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-229881>

Book Section

Submitted Version



The following work is licensed under a Creative Commons: Attribution 4.0 International (CC BY 4.0) License.

Originally published at:

Zhukova, Anastasia; Hamborg, Felix; Donnay, Karsten; Gipp, Bela (2022). XCoref: Cross-document Coreference Resolution in the Wild. In: Smits, Malte. Information for a Better World: Shaping the Global Future : 17th International Conference, iConference 2022, Virtual Event, February 28 – March 4, 2022, Proceedings, Part I. Cham: Springer, 272-291.

DOI: https://doi.org/10.1007/978-3-030-96957-8_25

XCoref: Cross-document Coreference Resolution in the Wild

Anastasia Zhukova¹[0000-0001-9084-2890], Felix Hamborg^{2,4}[0000-0003-2444-8056],
Karsten Donnay^{3,4}[0000-0002-9080-6539], and Bela Gipp^{1,4}[0000-0001-6522-3019]

¹ Data and Knowledge Engineering, University of Wuppertal, Germany
{last}@uni-wuppertal.de

² Dept. of Computer Science, University of Konstanz, Germany
felix.hamborg@uni-konstanz.de

³ Dept. of Political Science, University of Zurich, Switzerland donnay@ipz.uzh.ch

⁴ Heidelberg Academy of Sciences and Humanities, Germany

Abstract. Datasets and methods for cross-document coreference resolution (CDCR) focus on events or entities with strict coreference relations. They lack, however, annotating and resolving coreference mentions with more abstract or loose relations that may occur when news articles report about controversial and polarized events. Bridging and loose coreference relations trigger associations that may lead to exposing news readers to bias by word choice and labeling. For example, coreferential mentions of “direct talks between U.S. President Donald Trump and Kim” such as “an extraordinary meeting following months of heated rhetoric” or “great chance to solve a world problem” form a more positive perception of this event. A step towards bringing awareness of bias by word choice and labeling is the reliable resolution of coreferences with high lexical diversity. We propose an unsupervised method named *XCoref*, which is a CDCR method that capably resolves not only previously prevalent entities, such as persons, e.g., “Donald Trump,” but also abstractly defined concepts, such as groups of persons, “caravan of immigrants,” events and actions, e.g., “marching to the U.S. border.” In an extensive evaluation, we compare the proposed XCoref to a state-of-the-art CDCR method and a previous method TCA that resolves such complex coreference relations and find that XCoref outperforms these methods. Outperforming an established CDCR model shows that the new CDCR models need to be evaluated on semantically complex mentions with more loose coreference relations to indicate their applicability of models to resolve mentions in the “wild” of political news articles.

Keywords: cross-document coreference resolution · news analysis · media bias

1 Introduction

Coreference resolution (CR) is a set of techniques that aim to resolve mentions of entities, often in a single text document. CR is employed as an essential analysis component in a broad spectrum of use cases, e.g., to identify potential targets in sentiment analysis or as a part of discourse interpretation. While CR focuses on single documents, *cross-document coreference resolution* (CDCR) resolves concept mentions across a set of multiple documents. Compared to CR, CDCR is a less-researched task, although is

more challenging due to a larger search space and much required to facilitate content understanding of multiple articles. Further, many use cases require CDCR of more varying concepts than just entities typically resolved within CR, e.g., events or abstract entities.

Although the CDCR research has been gaining attention, the annotation schemes and corresponding datasets have been infrequently exploring a mix of identity and loose coreference relations and lexical diversity of the annotated chains of mentions. We explore CDCR in a particularly challenging use case, i.e., to identify bias by word choice and labeling. Such bias occurs due to strong variance in the words and loose coreference relations that yield possibly biased interpretations of events or entities. For example, coreferential mentions of “direct talks between U.S. President Donald Trump and Kim” such as “an extraordinary meeting following months of heated rhetoric” or “great chance to solve a world problem” form a more positive perception of this event.

Resolution of identity relations, i.e., coreference resolution, and resolution of more loose relations, i.e., bridging, are typically split into two separate tasks [17]. Resolution and evaluation of entity and event mentions of the two types of relation remains a research gap in general CDCR research. Hamborg et al. [12] first explored CDCR in a particularly challenging use case, i.e., to identify bias by word choice and labeling in news articles. Their proposed approach called TCA resolved mentions with strong lexical diversity.

In this paper, first, we revisit TCA and propose XCoref, an unsupervised sieve-based method that resolves jointly mentions of strict and loose identity relations into coreferential chains⁵. Methods of such design have been successfully used to resolve mentions of identity relation in events and entities [18, 22], and bridging relations in entities [14].

Second, we conduct an extensive evaluation where we compare XCoref with TCA and one of the state-of-the-art method for CDCR [2]. We evaluate the annotated mentions of varying coreference strength jointly as one task of set identification and calculate the results with the standard CoNLL metrics in (CD)CR [38]. We discuss a direction of the CDCR evaluation to address the complexity of coreferential chains annotated on diverse political news articles, i.e., in the “wild.”

2 Related work

Coreference resolution (CR) and cross-document coreference resolution (CDCR) are tasks that aim to resolve coreferential mentions in one or multiple documents, respectively [43]. (CD)CR approaches tend to depend on the annotation schemes of the CDCR datasets that specify the definition of mentions and coreferential relations [4].

Most (CD)CR datasets contain only strict identity relations, e.g., TAC KBP [27, 28], ACE [1, 20], MEANTIME [26], OntoNotes [45], ECB+ [3, 9]. Less commonly used (CD)CR datasets explore relations beyond strict identity. For example, NiDENT [40] is a CDCR dataset of entities-only mentions that was created by reannotating NP4E. NiDENT explores coreferential mentions of more loose coreference relations coined

⁵ For the peer review, an anonymous link to the code: <https://drive.google.com/file/d/18vfxdSwXE9fsjEq6IYilj9cDTJND0oXX/>

near-identity that among all included metonymy, e.g., “White House” to refer to the US government, and meronymy, e.g., “US president” being a part of the US government and representing it. Reacher Event Description (RED), a dataset for CR, contains also more loose coreference relations among events [32].

Mentions coreferential with more loose relations are harder to annotate and automatically resolve than mentions with identity relations [39]. Bridging relations occur when a connection between mentions is implied but is not strict, e.g., a “part-of” relation. Bridging relations, unlike identity relations, form a link between nouns that do not match in grammatical constraints, e.g., gender and number agreement, and allow linking noun and verb mentions, thus, constructing abstract entities [17]. The existing datasets for identification of bridging relations, e.g., ISNotes [14], BASHI [41], ARRAU [36], annotate the relations only of noun phrases on a single-document level and solve the problem as antecedent identification problem rather than identification of a set of coreferential anaphora [14]. Definition identification in DEFT dataset [44] focuses on annotating mentions that are linked with “definition-like” verb phrases (e.g., means, is, defines, etc.) but does not address linking the antecedents and definitions into the coreferential chains.

To our knowledge, only one dataset contains annotations of coreferential mentions with varying strength of coreferential relations. NewsWCL50 [12] contains annotations of concepts based on a minimum number of (coreferential) mentions across a set of news articles reporting on the same event. The dataset contains diverse concept types, such as actors, entities, events, geo-political entities (GPEs), and also more complex types, such as actions or abstract entities. The dataset argues that in political news articles more loose coreferential relations form links and associations to phrases that could bear bias by word choice and labeling, e.g., “DACA recipients” – “undocumented immigrants who came to the U.S. as children” – “illegal aliens” – “innocent kids.”

There are two approaches for event CDCR, easy-first and mention-pair [22]. Usually, easy-first approaches are unsupervised, whereas mention-pair are supervised. Most methods employ the easy-first approach and sequentially execute so-called sieves. Each sieve resolves mentions concerning specific characteristics. Earlier sieves target simple and generally reliable properties, such as heads of phrases. Later sieves address more difficult or specialized cases and use special techniques, such as pair-wise scoring of the pre-identified concepts with binary classifiers, e.g., SVM [12, 19, 21, 31]. Alternatively, a mention-pair approach uses a neural model trained to score the likelihood of a pair of the event- or entity-mentions to refer to the same semantic concept. The features to represent mentions are spans of text, contexts, and semantic dependencies [2].

Most CDCR methods focus on only events and resolve entities—if at all—as subordinate attributes of the events [2, 7, 10, 15, 16, 21, 22]. These a few CDCR methods resolve chains of both event and entity mentions with strict identity and there is only one method the resolve concepts with more loose the identity and coreference relations [12]. The method named TCA [12], resolves mentions of varying identity levels, i.e., include strict and near-identity. Target Concept Analysis (TCA) is a sieve-based method that resolves concepts that represent entities, events, and aggregating categories, e.g., categories that include mentions referring to both a country and its governmental institutions to which they belong annotated with a name of a country [12]. Other sieve-based feature-engineering methods for mention resolution were successfully used for identification

of entities and events with strict identity [18, 22, 31], and resolution of mentions with bridging relations [13, 14].

In conclusion, prior methods for CDCR suffer from at least one of two shortcomings, i.e., they (1) only resolve mentions interlinked with identity relations or (2) focus on event-driven narrowly defined coreferential mentions. The contributions of this paper are two-fold: first, we revisit the methodology of TCA because it is the only method that addresses the resolution of mentions with various identity relations [12]. Then, we propose XCoref, an unsupervised method that jointly resolves mentions with strict and loose anaphoric relations. Second, we evaluate the approach on a CDCR dataset with coreferential chains with varying identity relations, i.e., NewsWCL50, and compare the results to the previously proposed CDCR methods for these datasets using metrics established in the literature on (CD)CR, i.e., B3, CEAF_e, and MUC [29].

3 Methodology: XCoref

XCoref revisits Target Concept Analysis (TCA) proposed by Hamborg et al. [12]. XCoref consists of five sieves (see Figure 1) and applies the “easy-first” principle, i.e., it first resolves mentions that belong to named entities (NEs), such as person, organization, and country, and are coreferential with identity relation [14]. Afterward, the method addresses chains coreferential with mixed identity and bridging relations, i.e., groups of persons, events, and abstract entities. TCA is a sieve-based method to resolve coreferential chains, by addressing issues of the previously unresolved mentions. In contrast to TCA, XCoref resolves mentions of specific concept types in each sieve and analyzes combinations of phrases’ modifiers to resolve mentions with varying coreference relations.

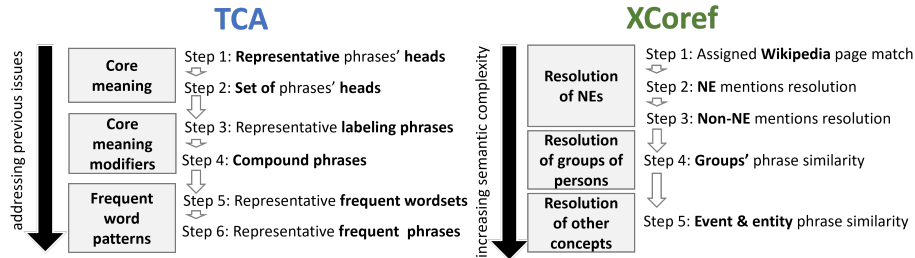


Fig. 1: Comparison of TCA to XCoref: TCA resolves mentions by addressing previous issues whereas XCoref first resolves coreferential chains of NEs with identity relations (S1-S3) and then resolve those chains with more loose coreferential relations (S4 and S5).

3.1 Preprocessing

Since each sieve of XCoref resolves specific concept types, we first need to identify these types. We distinguish nine types: person, group, country, misc, and their NE- and

non-NE variations [12]. Further, each sieve S_i uses a comparison matrix cm_i , i.e., a manually created cross-type matrix, to determine which types x and y to compare for potential merging ($cm_{i,x,y} \geq 1$ allows comparison, 0 not).

To yield reliable concept type determination, we propose a type scoring based on WordNet sense ranking [11]. For a given concept, Hamborg et al. [12] proposed counting frequencies of WordNet senses assigned to all concept’s heads’, e.g., “president” is four times “noun.person.” We propose an improvement to the counting by weighting the senses according to their rank in WordNet. Since highly ranked senses have more influence on the concept type score, the weighted scoring minimizes tilting of the final concept’s type towards the rare meanings of its mentions.

An additional improvement to the preprocessing proposed by Hamborg et al. [12] is CoreNLP’s coreference resolution on (virtually) concatenated documents. We observe that CoreNLP resolves more true mentions in concatenated texts than within single texts but is also prone to wrongly merging large coreferential chains, even when using the improved CR model [8]. Thus, we split chains by analyzing NEL results of each of their mentions, i.e., for each mention and its compound+head sub-phrases, we obtain its Wikipedia page name [23]. We split chains whose members are assigned to different Wikipages [26]. We attach a resolved Wikipedia title to each cleared chain as a new property.

Each mention contains dependency and structure parsing subtrees. Structure subtrees are created by 1) parsing the sentences of mentions’ origin, 2) mapping heads of mentions to the structure parsing trees, 3) taking the longest subtree but not larger than 20 tokens. The dependency subtree contains all the tokens as in the structure subtree. These subtrees play a role of the feature sources for the sieves.

S1: named entity linking (NEL)

The first sieve, S1, leverages CoreNLP to mainly resolve mentions NE-containing chains, i.e., pronominal, nominal, and pronoun mentions. Specifically, S1 reused CoreNLP’s coreference resolution on the concatenated documents and the assigned Wikipages. We resolve mentions by the winner-takes-it principle [14], i.e., we merge the smaller chains into the coreference chains that were pre-identified by CoreNLP if the chains have identical Wikipedia titles.

S2: head-word and compound match of NEs

The second sieve, S2, merges the NE-containing chains, which mentions CoreNLP and NEL (in S1) failed to resolve. These typically smaller NE-containing chains of comparable types by cm_2 are merged into the larger NE-containing chains if they have identical NE heads (“*Kim*” – “North Korean dictator *Kim Jong Un*”) or have cross-overlapping NE heads and compounds (“*Donald*” – “*Donald Trump*”).

S3: resolution of non-NE mentions

The third sieve, S3, merges any comparable NE-containing chain c_{ne} with a smaller chain with non-NE mentions c_{nn} if they match due to string or cosine similarity, e.g.,

“Teresa May” – “the prime minister.” S3 focuses on using dependency subtrees of the mentions and extracts head modifiers from the subtrees, e.g., adjective, compound, apposition, or noun. Such anaphora tend to be missed by CoreNLP. In S3, we merge chains if at least one of the conditions holds:

- 1) $r(c_{nn}) \subset r(c_{ne}) \wedge m(c_{nn}) \subset |r(c_{nn}) \cap r(c_{ne})|$, where $r(\cdot)$ extracts representative phrases, i.e., phrases that contain heads of phrases and their direct modifiers (i.e., adjectival, noun, and compound) expanded with a list of all apposition modifiers, and $m(\cdot)$ are modifiers (i.e., compounds and appositions) extracted from the chains’ mentions,
- 2) $|c_{nn} \cap c_{ne}| \geq 2$ and at least one head of phrases from each chain belongs to this intersection,
- 3) $\cos(v(c_{ne}), v(c_{nn})) \geq t_{nn}$, where $v(\cdot)$ is a mean vector of the word vector representation of all unique non-stop-words from the concept’s mentions, \cos is cosine similarity and t_{nn} is a threshold derived during experiments.

S4: identification of groups of persons

The fourth sieve, S4, resolves chains of mentions referring to the group of individuals, e.g., “illegal aliens” – “undocumented immigrants,” and semantically related phrases to countries, e.g., “Trump administration officials” – “American government.”

We incorporate a clustering approach of Zhukova et al. [46] that clusters mentions into chains in the decreasing semantic similarity between mentions. First, it finds distinctive cluster cores, i.e., each core is a group of mentions that are highly semantically similar to each other as to two factors: 1) cosine similarity of word-vectorized mentions in a potential core and 2) the normalized number of other mentions to which mentions of a potential core are similar. Then, the approach tries to assign the remaining mentions to the cores based on the relaxed rules of one-level similarity, i.e., using only cosine similarity between unresolved mentions and the already clustered mentions into chains. Mentions are to be left unresolved and converted into singleton-concepts if they do not meet similarity requirements, such as similarity to one of the core points.

Lastly, because the annotation of GPEs include also mentions to the nations of the countries, we merge the country-NE-containing chains c_{ne} with the groups of people c_{gr} if $\cos(v(h(c_{ne}), v(m(c_{gr}))) \geq t_{gr}$, where $h(\cdot)$ extracts heads of phrases, $m(\cdot)$ extracts modifiers (i.e., NE-compounds and NE-adjectival) from chains’ mentions, and t_{gr} is a threshold derived during experiments.

S5: events and abstract entities

The fifth sieve, S5, resolves mention chains of events and so-called abstract entities [36], i.e., actions, objects, events, etc, that contain mentions of noun phrases (NPs) and verb phrases (VPs), for example, “Trump-Kim *meeting*” – “*discussed* an issue”. Such mentions are typically not resolved by CoreNLP and contain either one the mentions or identical mentions. S5, first, keeps the unique mentions within chains, removes articles, and lemmatizes words. Second, S5 adapts a weighting scheme of word embedding vectors from Zhukova et al. [46]. S5 weights heads of mentions with a coefficient $k = 2$ and keeps the original vectors for the other words in a mention. Such a weighting scheme allows emphasising the core meaning of a mention while keeping mention’s

context. Lastly, we form coreferential chains by clustering the mentions with hierarchical clustering [30] using cosine distance, average linkage, and a threshold t_{cl} . Hierarchical clustering commonly used in concept identification [6].

4 Evaluation

We compare XCoref to the established CDCR model of Barhom et al. [2] on NewsWCL50, i.e., a concept-driven CDCR dataset with a mix of identify and bridging relations. We evaluate the datasets containing coreference anaphora with bridging relations with the standard coreference metrics of CoNLL [38], although most of the bridging relations are not classified as a problem of set identification [14].

4.1 Dataset: NewsWCL50

Similar to [7], we do not evaluate the approaches on two separate lists of mentions for events and entities, but consider the mentions of all mentions combined as if they formed abstract entities. If an approach requires input or separate event and entity mentions (e.g., Barhom et al. [2]), then we implement splitting of mentions into two lists based on the heads of phrases, i.e., VPs represent events and other part-of-speech tags represent entities.

We removed “ACTOR-I” category from NewsWCL50 due to a large level of abstractness [12]. Our manual inspection of the chains labeled with this category did not identify consistency in annotating mentions as coreferential by identity, near-identity, or bridging relations. Therefore, to focus only on the listed relations, we removed this category from the evaluation.

4.2 Methods and baselines

We compare the performance of XCoref to a lemma baseline, an established event-entity CDCR model of Barhom et al. [2], and Target Concept Analysis method (TCA) of Hamborg et al. [12] to identify semantic concepts in NewsWCL50. Additionally, we perform an ablation study and evaluate modifications of XCoref. For all approaches, we use the same default parameters across topics, datasets, and run configurations to facilitate fair evaluation. We describe each method briefly in the following.

Lemma Our baseline is a primitive CDCR approach that resolves mentions based on matching lemmas of the phrases’ heads. This baseline was also used by Barhom et al. [2] and establishes a fair comparison to the other approaches in the evaluation.

EeCDCR Barhom et al. [2] proposed a joint event and entity CDCR model (hereafter EeCDCR). EeCDCR is trained on ECB+ and resolves event and entity mentions jointly. To reproduce EeCDCR’s performance, we use the model’s full set of optional features: semantic role labeling (SRL), ELMo word embeddings [35], and dependency relations. Barhom et al. [2] used the output files of the SRL parser, SwiRL, which makes it

impossible to apply EeCDRCR to the other datasets. Therefore, we used the most recent AllenNLP’s SRL method [42] to make the feature extraction a part of the EeCDRCR and applicable for all datasets. To resolve intra-document mentions, identical to the original setup, we use Stanford CoreNLP. We reused default parameters for the model inference.

TCA Target Concept Analysis (TCA) is a method of identification of the reported concepts in the related news articles, the mentions of which are typically a subject of bias by word choice and labeling [12]. Based on the functionality of TCA, we can classify the method as CDCR focusing on the resolution of concepts with varying strength of cross-mention identity relations and mentions of high lexical diversity.

TCA uses six sieves to determine whether two candidate chains should be merged because they refer to the same semantic concept. Each sieve uses specific similarity measures, e.g., cosine similarity or string equality, and analyzes specific characteristics of the candidates’ mentions, e.g., heads and their modifiers, i.e., adjective, noun, apposition, compound, and number. We use the reported TCA’s default parameters for all datasets. For a word embedding model, we used a version of word2vec [25] that unlike the original implementation vectorizes out-of-vocabulary words [33].

Identical to Hamborg et al., we use TCA’s default parameters for all datasets. We report the results of multiple variants of TCA. First, the original version as reported by Hamborg et al. Second, to facilitate comparability of TCA’s sieves to the sieves of XCoref, we use the preprocessing steps of XCoref in TCA and also show the effectiveness of these steps to improve the performance. Additionally, we evaluate TCA_{preproc} using three different word embeddings (see below).

Ablation study For the ablation study, we test two variants of XCoref with which we investigate the effectiveness of various word embedding models and the effectiveness of approaches identifying more loose anaphoric coreference relations.

First, an “intermediate” model named $XCoref_{\text{interm}}$ uses sieves S1-S3 of XCoref, $S4_{\text{interm}}$ is a baseline used by Zhukova et al. [46] with the same threshold parameters, and $S5_{\text{interm}}$ is the second sieve adopted from TCA, i.e., cosine similarity of phrases’ heads. Resolution with semantically similar heads can effectively resolve mentions of abstract entities, e.g., “meeting” – “talks.” $S4_{\text{interm}}$ and $S5_{\text{interm}}$ resolve mentions of the same concept types as in XCoref. Using $XCoref_{\text{interm}}$, we test if the proposed methods for resolution of the bridging coreference anaphora outperform the simpler methods for the same coreference relations.

Second, we test for either XCoref and $XCoref_{\text{interm}}$ how using state-of-the-art, non-contextualized word embeddings affects their performance: word2vec [25], fastText [24], and GloVe [34]. We use the model implementations that facilitate the representation of out-of-vocabulary (OOV) words [33], which is critical to address the inability of the default word2vec and GloVe models to represent OOV words.

4.3 Metrics

We report the established CoNLL metrics for (CD)CR, i.e., MUC , $B3$, $CEAF_e$, and an average of them as $F1_{CoNLL}$ [38]. Similar to Barhom et al. [2], we evaluate the

methods with an official CoNLL scorer [37]. We do not distinguish between the strength of coreferential relations and evaluate the coreferential chains as if their mentions had relations of identical strength.

Method	Word vectors	MUC			B3			CEAF _e			F1 _{CoNLL}
		R	P	F1	R	P	F1	R	P	F1	
Lemma	—	75.7	93.8	83.8	36.8	88.8	52.0	63.1	7.8	13.9	49.9
EeCDCR	GloVe	69.4	90.6	78.6	33.1	82.3	47.2	58.6	7.8	13.7	46.5
TCA	word2vec	73.4	89.4	80.6	37.2	73.7	49.4	51.9	8.7	14.9	48.3
TCA _{preproc}	word2vec	72.9	89.5	80.3	38.4	75.6	50.9	54.5	9.1	15.6	48.9
	fastText	72.9	87.6	79.6	37.3	71.5	49.0	52.0	9.5	16.0	48.2
	GloVe	77.2	88.3	82.4	41.8	67.0	51.4	52.9	12.1	19.6	51.2
XCoref _{interm}	word2vec	68.4	90.3	77.8	37.7	84.0	52.0	63.0	8.4	14.8	48.2
	fastText	74.2	87.3	80.2	38.7	71.5	50.2	58.4	11.6	19.4	50.0
	GloVe	75.7	88.5	81.6	40.6	72.1	52.0	58.9	12.1	20.0	51.2
XCoref	word2vec	70.7	89.8	79.1	36.3	82.4	50.4	63.0	9.4	16.3	48.6
	fastText	78.6	90.0	83.9	43.1	70.5	53.5	60.4	13.7	22.4	53.3
	GloVe	79.3	90.8	84.7	44.4	72.2	55.0	61.1	13.9	22.6	54.1

Table 1: Evaluation of a lemma baseline, EeCDCR [2], TCA [12], XCoref_{interm} (a version of XCoref with baseline methods for sieves S4 and S5), and XCoref on the two diverse CDCR datasets: ECB+ and NewsWCL50.

Sieves	MUC			B3			CEAF _e			F1 _{CoNLL}
	R	P	F1	R	P	F1	R	P	F1	
init _{shared}	28.6	89.8	43.4	15.6	96.0	26.8	41.4	2.1	4.1	24.7
S1 _{shared}	40.2	92.3	56.0	22.9	95.0	36.9	49.7	3.1	5.8	32.9
S2 _{shared}	42.2	92.5	58.0	24.6	94.6	39.1	51.2	3.3	6.2	34.4
S3 _{shared}	45.7	91.0	60.8	27.3	91.3	42.1	51.6	3.6	6.7	36.5
S4 _{interm}	52.7	90.2	66.6	29.1	87.5	43.6	51.5	4.2	7.8	39.3
S5 _{interm}	75.7	88.5	81.6	40.6	72.1	52.0	58.9	12.1	20.0	51.2
S4	54.6	91.1	68.3	30.4	86.1	44.9	51.9	4.4	8.1	40.4
S5	79.3	90.8	84.7	44.4	72.2	55.0	61.3	13.9	22.7	54.1

Table 2: Comparison of XCoref’s sieves to the intermediate version of XCoref_{interm} in NewsWCL50 with GloVe word vectors. XCoref and XCoref_{interm} share sieves S1-S3 and differ in the last sieves. S4 and S5 of XCoref outperform sieves of XCoref_{interm}.

4.4 Results and Discussion

The evaluation shows that XCoref, our multi-sieve feature-based unsupervised method, is capable of consistently resolving coreferential anaphora of mixed identity and bridging coreference relations with $F1_{CoNLL} = 54.1$ when using GloVe as word vectors.

Method	Word vectors	MUC			B3			CEAF _e			F1 _{CoNLL}
		R	P	F1	R	P	F1	R	P	F1	
Lemma	—	74.8	71.7	73.2	17.4	68.8	27.8	53.8	5.3	9.7	36.9
EeCDCR	GloVe	67.2	54.9	60.4	16.7	56.0	25.7	49.2	4.4	8.0	31.4
TCA	word2vec	74.7	48.4	58.7	21.2	39.6	27.6	40.9	5.4	9.5	32.0
S3 _{interm}	word2vec	60.0	75.5	66.9	13.8	75.0	23.3	46.1	3.5	6.5	32.2
	fastText	71.6	68.7	70.1	21.8	60.7	32.1	38.8	4.4	8.0	36.7
	GloVe	72.6	72.1	72.3	19.8	63.0	30.1	38.1	4.5	8.1	36.8
S3	word2vec	80.0	69.2	74.2	27.3	55.0	36.5	45.7	7.2	12.5	41.1
	fastText	81.5	60.1	69.2	33.2	42.7	37.4	34.1	6.1	10.3	39.0
	GloVe	81.7	63.3	71.3	28.4	49.7	36.2	38.0	6.7	11.4	39.6

Table 3: Comparison of S4 for resolution groups of persons between XCoref to the intermediate version of XCoref_{interm}.

Section 4.3 reports the $F1_{CoNLL}$ score typically used to evaluate (CD)CR approaches (named $F1$ in this section). We find that XCoref outperforms all methods on NewsWCL50 dataset ($F1 = 54.1$). XCoref performs better than a simple lemma-baseline by $F1_{\Delta} = 4.2$ and outperforms TCA, i.e., a previously used method for the concept identification in NewsWCL50 dataset, by $F1_{\Delta} = 5.8$. XCoref shows that the mentions of the mixed identity and bridging relations are resolved if treating all mentions without an upfront separation by the coreference strength of the anaphora. Moreover, learning context can substituted by feature engineering from the phrases’ extracted nearest context, i.e., parsing subtrees.

When comparing TCA to its modification TCA_{preproc}, we see that the preprocessing steps employed in XCoref improve the overall performance, i.e., the improved identification of concept type and initial coreference resolution by CoreNLP on the combined documents instead on single documents. Using TCA with GloVe wordvectors improves the performance of original TCA by $F1_{\Delta} = 2.9$. We assume that the concentration of the NE mentions in NewsWCL50 is high and such preprocessing step positively impacts the overall performance.

Section 4.3 compares the sieves of XCoref to the “intermediate” version XCoref_{interm}. Sieves S4 and S5 outperform S4_{interm} and S5_{interm} by $\Delta_{S4} = 1.1$ and $\Delta_{S5} = 1.8$ correspondingly. Additionally, we evaluate S4 on the mentions annotated as groups of persons. Section 4.3 shows that S4 in the original implementation by Zhukova et al. [46] with word2vec as word embedding performs best and the implementation with GloVe is second best. We assume that the thresholds of the proposed approach are better suited for word2vec and could be tuned for GloVe.

Section 4.3 shows multiple examples of how XCoref manages to resolve the annotated mentions that belong to the complex abstract entities. The table shows that some mentions form coreference relations only within a given context. For example, a mention “a great chance to solve world problem,” which describes the summit between Donald Trump and Kim Jong Un, forms a positive association “the summit = a great chance” via a context-specific bridging relation. A coreference link “a one-to-one meeting with

Name	Resolved mentions
PRK-USA Summit	the summit meeting, a potential meeting of the two leaders, an extraordinary meeting following months of heated rhetoric, meet with the North Korean dictator, discuss its nuclear weapons program, Kim’s offer for a summit, a great chance to solve a world problem, won’t even have a meeting at all, a once-unthinkable encounter between him and Mr. Kim, a one-on-one meeting with North Korea leader Kim Jong Un, direct talks between U.S.President Donald Trump and Kim, Mr. Kim’s invitation to meet, the upcoming summit meeting with the North Korean leader, a great chance to solve a world problem, unwavering determination in addressing the challenge of North Korea
DNC’s lawsuit	the process of legal discovery, a sham lawsuit about a bogus Russian collusion claim, a bogus Russian collusion claim, allegations of obstruction of justice, a desperate attempt to keep a collusion narrative going ahead of November mid-term elections, a new law to raise money, the DNC’s move, the lawsuit to drum up donations for the party
Denuclearization	a deal to destroy only inter-continental missiles that could reach the United States, gives up nuclear weapons, months of heated rhetoric over Pyongyang’s nuclear weapons program, to engage in a process headed toward an ambiguous goal, broad and ”abstract ” statements about the need for North Korea to “denuclearize”, give up its nuclear program, yet to take any tangible steps to give up its nuclear arsenal, to address the threats posed by its nuclear and missile program
Coming into the US	made their way to the U.S.-Mexico border, begun crossing into the U.S., the arrival of a caravan of Central American migrants, the arrival of a few hundred, crossing through a legal port of entry, driving families to flee, the caravan’s steady approach to the U.S., to pass through into our country, several attempted illegal entries by people associated with the caravan, may be detained or fitted with ankle monitors and released, wait to be processed by U.S. authorities, to turn themselves in, a process that unfolds over several months or longer, made their way toward the border, a test of President Trump’s anti-immigrant politics, a volatile flash point in the immigration debate ignited by Mr.Trump, headed north together as a form of protection, the final act of the caravan
Immigrants	no legitimate asylum-seekers, the asylum-seekers, the individual migrants planning to seek asylum, groups of the migrants with their children, migrant families that request asylum, unauthorized immigrants, the migrants in the caravan, a caravan of immigrants, members of the caravan, these large “Caravans” of people, a few hundred asylum seekers, refugees, those individuals, applicants, people who request protection, people traveling without documents, several groups of people associated with the caravan, undocumented immigrants, asylum-seeking immigrant “caravan”, group of about 100 people

Table 4: Examples of mentions resolved by XCoref. The table shows the annotated concept names and only unique resolved mentions.

North Korean leader Kim Jong Un” – “a once unthinkable encounter between Trump and Mr. Kim” indicates similar positive bias via the chosen word choice.

NewsWCL50 dataset shows that CDCR datasets with a mix of identity and bridging, i.e., looser, coreference relations represent a challenge to the established models, such as EeCDCR (performs worst among all baselines). A new direction in CDCR research proposed Caciularu et al. [5] by training a cross-document language model to enable CDCR models to understand the broad and narrow context of the mentions. Such deep context learning will improve the representation of phrases’ semantics and cross-phrase relations.

5 Conclusion

We propose XCoref, an unsupervised sieve-based method for cross-document coreference resolution (CDCR) that unlike the state-of-the-art CDCR methods resolves mentions of a mix of strict and loose coreference relations, e.g., “American steelmakers” – “shuttered plants and mills,” and “the United States” - “Trump Administration officials.” XCoref performs best among the evaluated approaches and also outperforms a previous approach for resolution of such complex coreferential chains by $F1_{CoNLL} = 5.8$. Further, a well-established CDCR model performs worse on NewsWCL50. Our findings suggest that CDCR models need to be tested on more diverse CDCR datasets that contain both strict identity and more loose bridging coreference relations. In political news articles, i.e., a challenging “wild” environment for the application of CDCR approaches, such relations might create the biased perceptions of reported entities and concepts. Therefore, resolution of mentions with context-specific coreference relations is a step towards bringing awareness of bias by word choice and labeling and increase the level of complexity of CDCR research.

References

1. Linguistic Data Consortium, et al.: ACE (automatic content extraction) english annotation guidelines for events. version 5.4. 3. ACE (2005)
2. Barhom, S., Shwartz, V., Eirew, A., Bugert, M., Reimers, N., Dagan, I.: Revisiting joint modeling of cross-document entity and event coreference resolution. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 4179–4189. Association for Computational Linguistics, Florence, Italy (Jul 2019). <https://doi.org/10.18653/v1/P19-1409>, <https://www.aclweb.org/anthology/P19-1409>
3. Bejan, C.A., Harabagiu, S.: Unsupervised event coreference resolution with rich linguistic features. In: Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics. pp. 1412–1422 (2010)
4. Bugert, M., Reimers, N., Gurevych, I.: Generalizing cross-document event coreference resolution across multiple corpora. Computational Linguistics pp. 1–43 (2021)
5. Caciularu, A., Cohan, A., Beltagy, I., Peters, M.E., Cattan, A., Dagan, I.: Cross-document language modeling. CoRR **abs/2101.00406** (2021), <https://arxiv.org/abs/2101.00406>
6. Cambria, E., Poria, S., Bajpai, R., Schuller, B.: SenticNet 4: A semantic resource for sentiment analysis based on conceptual primitives. In: Proceedings of COLING 2016, the 26th international conference on computational linguistics: Technical papers. pp. 2666–2677 (2016)

7. Cattan, A., Eirew, A., Stanovsky, G., Joshi, M., Dagan, I.: Cross-document coreference resolution over predicted mentions. *CoRR abs/2106.01210* (2021), <https://arxiv.org/abs/2106.01210>
8. Clark, K., Manning, C.D.: 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). pp. 643–653. Association for Computational Linguistics, Berlin, Germany (Aug 2016). <https://doi.org/10.18653/v1/P16-1061>, <https://www.aclweb.org/anthology/P16-1061>
9. Cybulska, A., Vossen, P.: Using a sledgehammer to crack a nut? lexical diversity and event coreference resolution. In: Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14). pp. 4545–4552. European Language Resources Association (ELRA), Reykjavik, Iceland (May 2014), http://www.lrec-conf.org/proceedings/lrec2014/pdf/840_{_}Paper.pdf
10. Cybulska, A., Vossen, P.: Translating granularity of event slots into features for event coreference resolution. In: Proceedings of the The 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation. pp. 1–10. Association for Computational Linguistics, Denver, Colorado (Jun 2015). <https://doi.org/10.3115/v1/W15-0801>, <https://www.aclweb.org/anthology/W15-0801>
11. Fellbaum, C.: Wordnet and wordnets. In: Brown, K. (ed.) *Encyclopedia of Language and Linguistics*. pp. 665–670. Elsevier, Oxford (2005)
12. Hamborg, F., Zhukova, A., Gipp, B.: Automated identification of media bias by word choice and labeling in news articles. In: Proceedings of the ACM/IEEE Joint Conference on Digital Libraries (JCDL) (2019). <https://doi.org/10.1109/JCDL.2019.00036>
13. Hou, Y., Markert, K., Strube, M.: Global inference for bridging anaphora resolution. In: Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. pp. 907–917. Association for Computational Linguistics, Atlanta, Georgia (Jun 2013), <https://aclanthology.org/N13-1111>
14. Hou, Y., Markert, K., Strube, M.: Unrestricted Bridging Resolution. *Computational Linguistics* **44**(2), 237–284 (06 2018). https://doi.org/10.1162/coli_a.00315, https://doi.org/10.1162/coli_a_00315
15. Kenyon-Dean, K., Cheung, J.C.K., Precup, D.: Resolving event coreference with supervised representation learning and clustering-oriented regularization. In: Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics. pp. 1–10. Association for Computational Linguistics, New Orleans, Louisiana (Jun 2018). <https://doi.org/10.18653/v1/S18-2001>, <https://www.aclweb.org/anthology/S18-2001>
16. Keshtkaran, A., Yuhaniz, S.S., Ibrahim, S.: An overview of cross-document coreference resolution. In: 2017 International Conference on Computer and Drone Applications (IConDA). pp. 43–48 (2017)
17. Kobayashi, H., Ng, V.: Bridging resolution: A survey of the state of the art. In: Proceedings of the 28th International Conference on Computational Linguistics. pp. 3708–3721. International Committee on Computational Linguistics, Barcelona, Spain (Online) (Dec 2020). <https://doi.org/10.18653/v1/2020.coling-main.331>, <https://aclanthology.org/2020.coling-main.331>
18. Lee, H., Peirsman, Y., Chang, A., Chambers, N., Surdeanu, M., Jurafsky, D.: Stanford’s multi-pass sieve coreference resolution system at the conll-2011 shared task. In: Proceedings of the 15th conference on computational natural language learning: Shared task. pp. 28–34. Association for Computational Linguistics (2011)
19. Lee, H., Recasens, M., Chang, A., Surdeanu, M., Jurafsky, D.: Joint entity and event coreference resolution across documents. In: Proceedings of the 2012 Joint Conference on Empirical

- Methods in Natural Language Processing and Computational Natural Language Learning. pp. 489–500. Association for Computational Linguistics, Jeju Island, Korea (Jul 2012), <https://www.aclweb.org/anthology/D12-1045>
20. Linguistic Data Consortium, et al.: ACE (automatic content extraction) english annotation guidelines for entities. Tech. rep., Technical report, Linguistic Data Consortium (2008)
 21. Lu, J., Ng, V.: Event coreference resolution with multi-pass sieves. In: Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16). pp. 3996–4003. European Language Resources Association (ELRA), Portorož, Slovenia (May 2016)
 22. Lu, J., Ng, V.: Event coreference resolution: A survey of two decades of research. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18. pp. 5479–5486. International Joint Conferences on Artificial Intelligence Organization (7 2018). <https://doi.org/10.24963/ijcai.2018/773>, <https://doi.org/10.24963/ijcai.2018/773>
 23. MediaWiki: MediaWiki — MediaWiki, the free wiki engine (2020), <https://www.mediawiki.org/w/index.php?title=MediaWiki&oldid=3770840>, [Online; accessed 15-May-2020]
 24. Mikolov, T., Grave, E., Bojanowski, P., Puhersch, C., Joulin, A.: Advances in pre-training distributed word representations. In: Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018) (2018)
 25. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems 26, pp. 3111–3119. Curran Associates, Inc. (2013)
 26. Minard, A.L., Speranza, M., Urizar, R., Altuna, B., van Erp, M., Schoen, A., van Son, C.: MEANTIME, the NewsReader multilingual event and time corpus. In: Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16). pp. 4417–4422. European Language Resources Association (ELRA), Portorož, Slovenia (May 2016), <https://www.aclweb.org/anthology/L16-1699>
 27. Mitamura, T., Liu, Z., Hovy, E.: Overview of TAC KBP 2015 event nugget track. In: TAC (2015)
 28. Mitamura, T., Liu, Z., Hovy, E.H.: Events detection, coreference and sequencing: What's next? overview of the TAC KBP 2017 event track. In: TAC (2017)
 29. Moosavi, N.S., Strube, M.: Which coreference evaluation metric do you trust? a proposal for a link-based entity aware metric. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 632–642. Association for Computational Linguistics, Berlin, Germany (Aug 2016). <https://doi.org/10.18653/v1/P16-1060>, <https://www.aclweb.org/anthology/P16-1060>
 30. Murtagh, F., Contreras, P.: Algorithms for hierarchical clustering: an overview. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 2(1), 86–97 (2012)
 31. NLP Architect, Intel AI Lab: Cross document co-reference (2018), http://nlp_architect.nervanasys.com/cross_doc_coref.html, [accessed 10-January-2021]
 32. O'Gorman, T., Wright-Bettner, K., Palmer, M.: Richer event description: Integrating event coreference with temporal, causal and bridging annotation. In: Proceedings of the 2nd Workshop on Computing News Storylines (CNS 2016). pp. 47–56. Association for Computational Linguistics, Austin, Texas (Nov 2016). <https://doi.org/10.18653/v1/W16-5706>, <https://www.aclweb.org/anthology/W16-5706>
 33. Patel, A., Sands, A., Callison-Burch, C., Apidianaki, M.: Magnitude: A fast, efficient universal vector embedding utility package. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. pp. 120–126 (2018)

34. Pennington, J., Socher, R., Manning, C.: GloVe: Global vectors for word representation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 1532–1543. Association for Computational Linguistics, Doha, Qatar (Oct 2014). <https://doi.org/10.3115/v1/D14-1162>, <https://www.aclweb.org/anthology/D14-1162>
35. Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., Zettlemoyer, L.: 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). pp. 2227–2237. Association for Computational Linguistics, New Orleans, Louisiana (Jun 2018). <https://doi.org/10.18653/v1/N18-1202>, <https://www.aclweb.org/anthology/N18-1202>
36. Poesio, M., Artstein, R.: Anaphoric annotation in the ARRAU corpus. In: Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC’08). European Language Resources Association (ELRA), Marrakech, Morocco (May 2008), http://www.lrec-conf.org/proceedings/lrec2008/pdf/297_paper.pdf
37. Pradhan, S., Luo, X., Recasens, M., Hovy, E., Ng, V., Strube, M.: Scoring coreference partitions of predicted mentions: A reference implementation. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). pp. 30–35. Association for Computational Linguistics, Baltimore, Maryland (2014). <https://doi.org/10.3115/v1/P14-2006>, <https://www.aclweb.org/anthology/P14-2006>
38. Pradhan, S., Moschitti, A., Xue, N., Uryupina, O., Zhang, Y.: CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes. In: Joint Conference on EMNLP and CoNLL - Shared Task. pp. 1–40. Association for Computational Linguistics, Jeju Island, Korea (Jul 2012), <https://www.aclweb.org/anthology/W12-4501>
39. Recasens, M., Hovy, E., Martí, M.A.: A typology of near-identity relations for coreference (NIDENT). In: Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10) (2010)
40. Recasens, M., Martí, M.A., Orasan, C.: Annotating near-identity from coreference disagreements. In: Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12). pp. 165–172. European Language Resources Association (ELRA), Istanbul, Turkey (May 2012), http://www.lrec-conf.org/proceedings/lrec2012/pdf/674_{_}Paper.pdf
41. Rösiger, I.: BASHI: A corpus of Wall Street Journal articles annotated with bridging links. In: Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018). European Language Resources Association (ELRA), Miyazaki, Japan (May 2018), <https://aclanthology.org/L18-1058>
42. Shi, P., Lin, J.: Simple BERT models for relation extraction and semantic role labeling. arXiv preprint [arXiv:1904.05255](https://arxiv.org/abs/1904.05255) (2019)
43. Singh, S., Bhattacharjee, K., Darbari, H., Verma, S.: Analyzing coreference tools for NLP application. *International Journal of Computer Sciences and Engineering* 7, 608–615 (5 2019). <https://doi.org/https://doi.org/10.26438/ijcse/v7i5.608615>
44. Spala, S., Miller, N.A., Yang, Y., Derroncourt, F., Dockhorn, C.: DEFT: A corpus for definition extraction in free- and semi-structured text. In: Proceedings of the 13th Linguistic Annotation Workshop. pp. 124–131. Association for Computational Linguistics, Florence, Italy (Aug 2019). <https://doi.org/10.18653/v1/W19-4015>, <https://www.aclweb.org/anthology/W19-4015>
45. Weischedel, R., Pradhan, S., Ramshaw, L., Palmer, M., Xue, N., Marcus, M., Taylor, A., Greenberg, C., Hovy, E., Belvin, R., et al.: OntoNotes release 4.0. LDC2011T03, Philadelphia, Penn.: Linguistic Data Consortium (2011)

46. Zhukova, A., Hamborg, F., Donnay, K., Gipp, B.: Concept identification of directly and indirectly related mentions referring to groups of persons. In: Diversity, Divergence, Dialogue: 16th International Conference, iConference 2021, Beijing, China, March 17–31, 2021, Proceedings, Part I 16. pp. 514–526. Springer International Publishing (2021)