

Cross-document Coreference Resolution over Predicted Mentions

Arie Cattan¹ Alon Eirew^{1,2} Gabriel Stanovsky³ Mandar Joshi⁴ Ido Dagan¹

¹Computer Science Department, Bar Ilan University

²Intel Labs, Israel ³The Hebrew University of Jerusalem

⁴Allen School of Computer Science & Engineering, University of Washington, Seattle, WA

arie.cattan@gmail.com alon.eirew@intel.com

gabis@cse.huji.ac.il mandar90@cs.washington.edu

dagan@cs.biu.ac.il

Abstract

Coreference resolution has been mostly investigated within a single document scope, showing impressive progress in recent years based on end-to-end models. However, the more challenging task of cross-document (CD) coreference resolution remained relatively under-explored, with the few recent models applied only to gold mentions. Here, we introduce the first end-to-end model for CD coreference resolution from raw text, which extends the prominent model for within-document coreference to the CD setting. Our model achieves competitive results for event and entity coreference resolution on gold mentions. More importantly, we set first baseline results, on the standard ECB+ dataset, for CD coreference resolution over predicted mentions. Further, our model is simpler and more efficient than recent CD coreference resolution systems, while not using any external resources.¹

1 Introduction

Cross-document (CD) coreference resolution consists of identifying textual mentions across multiple documents that refer to the same concept. For example, consider the following sentences from the ECB+ dataset (Cybulska and Vossen, 2014), where colors represent coreference clusters (for brevity, we omit some clusters):

1. *Thieves pulled off a two million euro jewellery heist in central Paris on Monday after smashing their car through the store's front window.*
2. *Four men drove a 4x4 through the front window of the store on Rue de Castiglione, before making off with the jewellery and watches.*

Despite its importance for downstream tasks, CD coreference resolution has been lagging behind the

impressive strides made in the scope of a single document (Lee et al., 2017; Joshi et al., 2019, 2020; Wu et al., 2020). Further, state-of-the-art models exhibit several shortcomings, such as operating on gold mentions or relying on external resources such as SRL or a paraphrase dataset (Shwartz et al., 2017), preventing them from being applied on realistic settings.

To address these limitations, we develop the first end-to-end CD coreference model building upon a prominent within-document (WD) coreference model (Lee et al., 2017) which we extend with recent advances in transformer-based encoders. We address the inherently non-linear nature of the CD setting by combining the WD coreference model with agglomerative clustering that was shown useful in CD models. Our model achieves competitive results on ECB+ over gold mentions and sets baseline results over predicted mentions. Our model is also simpler and substantially more efficient than existing CD coreference systems. Taken together, our work seeks to bridge the gap between WD and CD coreference, driving further research of the latter in realistic settings.

2 Background

Cross-document coreference Previous works on CD coreference resolution learn a pairwise scorer between mentions and use a clustering approach to form the coreference clusters (Cybulska and Vossen, 2015; Yang et al., 2015; Choubey and Huang, 2017; Kenyon-Dean et al., 2018; Bugert et al., 2020). Barhom et al. (2019) proposed to jointly learn entity and event coreference resolution, leveraging predicate-argument structures. Their model forms the coreference clusters incrementally, while alternating between event and entity coreference. Based on this work, Meged et al. (2020) improved results on event coreference by leverag-

¹<https://github.com/ariecattan/coref>




Figure 1: A high-level diagram of our model for cross-document coreference resolution. (1) Extract and score all possible spans, (2) keep top spans according to $s_m(i)$, (3) score all pairs $s(i, j)$, and (4) cluster spans using agglomerative clustering.

ing a paraphrase resource (Chirps; Schwartz et al., 2017) as distant supervision. Parallel to our work, recent approaches propose to fine-tune BERT on the pairwise coreference scorer (Zeng et al., 2020), where the state-of-the-art on ECB+ is achieved using a cross-document language model (CDLM) on pairs of full documents (Caciularu et al., 2021). Instead of applying BERT for all mentions pairs which is quadratically costly, our work separately encodes each (predicted) mention.

All above models suffer from several drawbacks. First, they use only gold mentions and treat entities and events separately.² Second, pairwise scores are recomputed after each merging step, which is resource and time consuming. Finally, they rely on additional resources, such as semantic role labeling, a within-document coreference resolver, and a paraphrase resource, which limits the applicability of these models in new domains and languages. In contrast, we use no such external resources.

Within-document coreference The *e2e-coref* WD coreference model (Lee et al., 2017) learns for each span i a distribution over its antecedents.

Considering all possible spans as potential mentions, the scoring function $s(i, j)$ between span i and j , where j appears before i , has three components: the two mention scores $s_m(\cdot)$ of spans i and j , and a pairwise antecedent score $s_a(i, j)$ for span j being an antecedent of span i .

Each span is represented with the concatenation of four vectors: the output representations of the span boundary (first and last) tokens ($x_{\text{FIRST}(i)}$, $x_{\text{LAST}(i)}$), an attention-weighted sum of token representations $\hat{x}_i$, and a feature vector $\phi(i)$. These span representations (g_i) are first fed into a

mention scorer $s_m(\cdot)$ to filter the λT (where T is the number of tokens) spans with the highest scores. Then, the model learns for each of these spans to optimize the marginal log-likelihood of its correct antecedents. The full description of the model is described below:

$$\begin{aligned} g_i &= [x_{\text{FIRST}(i)}, x_{\text{LAST}(i)}, \hat{x}_i, \phi(i)] \\ s_m(i) &= \text{FFNN}_m(g_i) \\ s_a(i, j) &= \text{FFNN}_a([g_i, g_j, g_i \circ g_j]) \\ s(i, j) &= s_m(i) + s_m(j) + s_a(i, j) \end{aligned}$$

3 Model

The overall structure of our model is shown in Figure 1. The major obstacle in applying the *e2e-coref* model directly in the CD setting is its reliance on *textual ordering* – it forms coreference chains by linking each mention to an antecedent span appearing before it in the document. This linear clustering method cannot be used in the multiple-document setting since there is no inherent ordering between the documents. Additionally, ECB+ (the main benchmark for CD coreference resolution) is relatively small compared to OntoNotes (Pradhan et al., 2012), making it hard to jointly optimize mention detection and coreference decision. These challenges have implications in all stages of model development, as elaborated below.

Pre-training To address the small scale of the dataset, we pre-train the mention scorer $s_m(\cdot)$ on the gold mention spans, as ECB+ includes singleton annotation. This enables generating good candidate spans from the first epoch, and as we show in Section 4.3, it substantially improves performance.

Training Instead of comparing a mention only to its previous spans in the text, our pairwise scorer $s_a(i, j)$ compares a mention to all other spans across all the documents. The positive instances for training consist of all the pairs of highest scoring

²Few works (Yang et al., 2015; Choubey and Huang, 2017) do use predicted mentions, by considering the intersection of predicted and gold mentions for evaluation, and thus not penalizing models for false positive mention identification. Moreover, they used a different version of ECB+ with known annotation errors, as noted in Barhom et al. (2019).

mention spans that belong to the same coreference cluster, while the negative examples are sampled (20x the number of positive pairs) from all other pairs. The overall score is then optimized using the binary cross-entropy loss as follows:

$$L = -\frac{1}{|N|} \sum_{x,z \in N} y \cdot \log(s(x, z))$$

where N corresponds to the set of mention-pairs (x, z) , and $y \in \{0, 1\}$ to a pair label. Full implementation details are described in Appendix A.1. Notice that the mention scorer $s_m(\cdot)$ is further trained in order to generate better candidates at each training step. When training and evaluating the model in experiments over gold mentions, we ignore the span mention scores, $s_m(\cdot)$, and the gold mention representations are directly fed into the pairwise scorer $s_a(i, j)$.

Inference At inference time, we score all spans; prune spans with lowest scores; score the pairs; and finally form the coreference clusters using an agglomerative clustering (using average-linking method) over these pairwise scores, following common practices in CD coreference resolution (Yang et al., 2015; Choubey and Huang, 2017; Kenyon-Dean et al., 2018; Barhom et al., 2019). Since the affinity scores $s(i, j)$ are also computed for mention pairs in different documents, the agglomerative clustering can effectively find cross-document coreference clusters.

4 Experiments

4.1 Experimental setup

Following most recent work, we conduct our experiments ECB+ (Cybulska and Vossen, 2014), which is the largest dataset that includes both WD and CD coreference annotation (see Appendix A.2). We use the document clustering of Barhom et al. (2019) for pre-processing and apply our coreference model separately on each predicted document cluster.

Following Barhom et al. (2019), we present the model’s performance on both event and entity coreference resolution. In addition, inspired by Lee et al. (2012), we train our model to perform event and entity coreference jointly, which we term “ALL”. This represents a useful scenario when we are interested in finding all the coreference links in a set of documents, without having to distinguish event and entity mentions. Addressing CD coreference with ALL is challenging because (1) the

search space is larger than when treating separately event and entity coreference and (2) models need to make subtle distinctions between event and entity mentions that are lexically similar but do not corefer. For example, the entity *voters* do not corefer with the event *voted*.

We apply RoBERTa_{LARGE} (Liu et al., 2019) to encode the documents. Long documents are split into non-overlapping segments of up to 512 word-piece tokens and are encoded independently (Joshi et al., 2019). Due to memory constraints, we freeze output representations from RoBERTa instead of fine-tuning all parameters. For all experiments, we use a single GeForce GTX 1080 Ti 12GB GPU. The training takes 2.5 hours for the most expensive setting (ALL on predicted mentions), while inference over the test set takes 11 minutes.

4.2 Results

Table 1 presents the combined within- and cross-document results of our model, in comparison to previous work on ECB+. We report the results using the standard evaluation metrics MUC, B³, CEAF, and the average F1 of these metrics, called CoNLL F1 (main evaluation).

When evaluated on gold mentions, our model achieves competitive results for event (81 F1) and entity (73.1) coreference. In addition, we set baseline results where the model does not distinguish between event and entity mentions at inference time (denoted as the ALL setting). The overall performance on ECB+ obtained using two separate models for event and entity is negligibly higher (+0.6 F1) than our single ALL model.

Our model is the first to enable end-to-end CD coreference on *raw* text (predicted mentions). As expected, the performance is lower than that using gold mentions (e.g 26.6 F1 drop in event coreference), indicating the large room for improvement over predicted mentions. It should be noted that beyond mention detection errors, two additional factors contribute to the performance drop when moving to predicted mentions. First, while WD coreference systems typically disregard singletons (mentions appearing only once) when evaluating on *raw* text, CD coreference models do consider singletons when evaluating on *gold* mentions on ECB+. We observe that this difference affects the evaluation, explaining about 10% absolute points out of the aforementioned drop of 26.6. The effect of singletons on coreference evaluation is further

		MUC			B ³			CEAF _{Fe}			CoNLL
		R	P	F ₁	R	P	F ₁	R	P	F ₁	F ₁
Event	Barhom et al. (2019)	77.6	84.5	80.9	76.1	85.1	80.3	81.0	73.8	77.3	79.5
	Meged et al. (2020)	78.8	84.7	81.6	75.9	85.9	80.6	81.1	74.8	77.8	80.0
	Our model – Gold	85.1	81.9	83.5	82.1	82.7	82.4	75.2	78.9	77.0	81.0
	Zeng et al. (2020)*	85.6	89.3	87.5	77.6	89.7	83.2	84.5	80.1	82.3	84.6
	Caciularu et al. (2021)*	87.1	89.2	88.1	84.9	87.9	86.4	83.3	81.2	82.2	85.6
	Our model – Predicted	66.6	65.3	65.9	56.4	50.0	53.0	47.8	41.3	44.3	54.4
Entity	Barhom et al. (2019)	78.6	80.9	79.7	65.5	76.4	70.5	65.4	61.3	63.3	71.2
	Our model – Gold	85.7	81.7	83.6	70.7	74.8	72.7	59.3	67.4	63.1	73.1
	Caciularu et al. (2021)*	88.1	91.8	89.9	82.5	81.7	82.1	81.2	72.9	76.8	82.9
	Our model – Predicted	43.5	53.1	47.9	25.0	38.9	30.4	31.4	26.5	28.8	35.7
ALL	Our model – Gold	84.2	81.6	82.9	76.8	77.5	77.1	68.4	72.4	70.3	76.7
	Our model – Predicted	49.7	58.5	53.7	33.2	46.5	38.7	40.4	35.2	37.6	43.4

Table 1: Combined within- and cross-document results on the ECB+ test set, for event, entity and the unified task, that we term ALL. Our results (in italics) are close to state-of-the-art (in bold) for event and entity over gold mentions, while they set a new benchmark result over predicted mentions and for the ALL setting. The run-time complexity in (Zeng et al., 2020; Caciularu et al., 2021) is substantially more expensive because they apply BERT and CDLM for every mention-pair with their corresponding context (sentence and full document).

	Gold		Predicted	
	WD	CD	WD	CD
Event	86.6	81.0	59.6	54.4
Entity	81.2	73.1	39.7	35.7
ALL	83.9	76.7	46.3	43.4

Table 2: Results (CoNLL F1) of our model, on within-document (WD) vs. cross-document (CD), using gold and predicted mentions. For all settings, results on WD are higher, indicating the need in addressing typical challenges of CD coreference resolution.

	Gold	Δ	Predicted	Δ
Event	76.0	-5.0	48.2	-6.2
Entity	70.9	-2.2	34.4	-1.3
ALL	74.1	-2.6	41.4	-2.0

Table 3: CoNLL F1 results of our model without document clustering, using gold and predicted mentions.

explored in (Cattan et al., 2021). Second, entities are annotated in ECB+ only if they participate in event, making participant detection an additional challenge. This explains the more important performance drop in entity and ALL.

Table 2 presents the CoNLL F1 results of within- and cross-document coreference resolution for both gold and predicted mentions on ECB+. For all settings, results are higher in within-document coreference resolution, showing the need in addressing typical challenges of CD coreference resolution.

Table 3 shows the results of our model without document clustering. Here, the performance drop and error reduction are substantially larger for event coreference (-6.2/12%) than entity coreference (-1.3/2%) and ALL (-2/3.5%). This difference is probably due to the structure of ECB+ which poses a lexical ambiguity challenge for events, while the document clustering step reconstructs almost perfectly the original subtopics, as shown in (Barhom et al., 2019).

Further, the higher results on event coreference do not mean that the task is inherently easier than entity coreference. In fact, when ignoring singletons in the evaluation, as done on OntoNotes, the performance of event coreference is lower than entity coreference (62.1 versus 65.3 CoNLL F1) (Cattan et al., 2021). This happens because event singletons are more common compared to entity singletons (30% vs. 17%), as shown in Appendix A.2.

Finally, our model is more efficient in both training and inference since the documents are encoded using just one pass of RoBERTa, and the pairwise scores are computed only once using a simple MLP. For comparison, previous models compute pairwise scores at each iteration (Barhom et al., 2019; Meged et al., 2020), or apply a BERT-model to every mention pairs with their sentence (Zeng et al., 2020) or full document (Caciularu et al., 2021).³

³For a rough estimation, our model runs for 2 minutes while Barhom et al. (2019)’s model runs for 37 minutes on similar hardware.

	F1	Δ
Our model	58.1	
– pre-train of mention scorer	54.9	–3.2
– dynamic pruning	54.1	–4.0
– negative sampling	56.7	–1.4

Table 4: Ablation results (CoNLL F1) of our model on the development set of ECB+ event coreference.

4.3 Ablations

To show the importance of each component of our model, we ablate several parts and compute F1 scores on the development set of the ECB+ event dataset. The results are presented in Table 4 using predicted mentions without document clustering.

Skipping the *pre-training* of the mention scorer results in a 3.2 F1 points drop in performance. Indeed, the relatively small training data in the ECB+ dataset (see Appendix A.2) might be not sufficient when using only end-to-end optimization, and pre-training of the mention scorer helps generate good candidate spans from the first epoch.

To analyze the effect of the dynamic pruning, we froze the mention scorer during the pairwise training, and kept the same candidate spans along the training. The significant performance drop (4 F1) reveals that the mention scorer inherently incorporates coreference signal.

Finally, using all negative pairs for training leads to a performance drop of 1.4 points and significantly increases the training time.

4.4 Qualitative Analysis

We sampled topics from the development set and manually analyzed the errors of the ALL configuration. The most common errors were due to an over-reliance on lexical similarity. For example, the event “*Maurice Cheeks was fired*” was wrongly predicted to be coreferent with a similar, but different event, “*the Sixers fired Jim O’Brien*”, probably because of related context, as both coached the Philadelphia 76ers. On the other hand, the model sometimes struggles to merge mentions that are lexically different but semantically similar (e.g. “*Jim O’Brien was shown the door*”, “*Philadelphia fire coach Jim O’Brien*”). The model also seems to struggle with temporal reasoning, in part due to missing information. For example, news articles from different days have different relative reference to time, while the publication date of the articles is not always available. As a result, the model missed

linking “*Today*” in one document to “*Saturday*” in another document.

5 Conclusion and Discussion

We developed the first end-to-end baseline for CD coreference resolution over predicted mentions. Our simple and efficient model achieve competitive results over gold mentions on both event and entity coreference, while setting baseline results for future models over predicted mentions.

Nonetheless, we note a few limitations of our model that could be addressed in future work. First, following most recent work on cross-document coreference resolution (§2), our model requires $\mathcal{O}(n^2)$ pairwise comparisons to form the coreference cluster. While our model is substantially more efficient than previous work (§4.2), applying it on a large-scale dataset would involve a scalability challenge. Future work may address the scalability issue by using recent approaches for hierarchical clustering on massive datasets (Monath et al., 2019, 2021). Another appealing approach consists of splitting the corpus into subsets of documents, constructing initial coreference clusters (in parallel) on the subsets, then merging meta-clusters from the different sets. We note though that it is currently impossible to test such solutions for more extensive scalability, pointing to a need in collecting larger-scale datasets for cross-document coreference. Second, to improve overall performance over predicted mentions, future work may incorporate, explicitly or implicitly, semantic role labeling signals in order to identify event participants for entity prediction, as well as for better event structure matching. Further, dedicated components may be developed for mention detection and coreference linking, which may be jointly optimized.

Acknowledgments

We thank the anonymous reviewers for their helpful comments. The work described herein was supported in part by grants from Intel Labs, Facebook, the Israel Science Foundation grant 1951/17, the Israeli Ministry of Science and Technology, the German Research Foundation through the German-Israeli Project Cooperation (DIP, grant DA 1600/1-1), and from the Allen Institute for AI.

Ethical Considerations

Dataset As mentioned in the paper, we use the ECB+ dataset, available at

<http://www.newsreader-project.eu/results/data/the-ecb-corpus/>.

Model Our cross-document coreference model does not contain any intentional biasing or ethical issues. As mentioned in the paper (§4.1), we conduct our experiments on a single 12GB GPU, and both training and inference times are relatively low.

References

- Shany Barhom, Vered Shwartz, Alon Eirew, Michael Bugert, Nils Reimers, and Ido Dagan. 2019. [Re-visiting joint modeling of cross-document entity and event coreference resolution](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4179–4189, Florence, Italy. Association for Computational Linguistics.
- Cosmin Bejan and Sanda Harabagiu. 2010. [Unsupervised event coreference resolution with rich linguistic features](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1412–1422, Uppsala, Sweden. Association for Computational Linguistics.
- Michael Bugert, N. Reimers, and Iryna Gurevych. 2020. Cross-document event coreference resolution beyond corpus-tailored systems. *ArXiv*, abs/2011.12249.
- Avi Caciularu, Arman Cohan, Iz Beltagy, Matthew E. Peters, Arie Cattan, and Ido Dagan. 2021. Cross-document language modeling. *ArXiv*, abs/2101.00406.
- Arie Cattan, Alon Eirew, Gabriel Stanovsky, Mandar Joshi, and Ido Dagan. 2021. Realistic evaluation principles for cross-document coreference resolution. In *Proceedings of the Joint Conference on Lexical and Computational Semantics*. Association for Computational Linguistics.
- Prafulla Kumar Choubey and Ruihong Huang. 2017. [Event coreference resolution by iteratively unfolding inter-dependencies among events](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2124–2133, Copenhagen, Denmark. Association for Computational Linguistics.
- Agata Cybulska and Piek Vossen. 2014. [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)*, pages 4545–4552, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Agata Cybulska and Piek Vossen. 2015. [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*, pages 1–10, Denver, Colorado. Association for Computational Linguistics.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256.
- 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.
- Kian Kenyon-Dean, Jackie Chi Kit Cheung, and Doina Precup. 2018. [Resolving event coreference with supervised representation learning and clustering-oriented regularization](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 1–10, New Orleans, Louisiana. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Heeyoung Lee, Marta Recasens, Angel Chang, Mihai Surdeanu, and Dan Jurafsky. 2012. [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*, pages 489–500, Jeju Island, Korea. Association for Computational Linguistics.
- 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.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Yehudit Meged, Avi Caciularu, Vered Shwartz, and Ido Dagan. 2020. [Paraphrasing vs coreferring: Two sides of the same coin](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*,

- pages 4897–4907, Online. Association for Computational Linguistics.
- Nicholas Monath, Kumar Avinava Dubey, Guru Guruganesh, M. Zaheer, Amr Ahmed, A. McCallum, Gokhan Mergen, Marc Najork, Mert Terzihan, B. Tjanaka, Yuan Wang, and Yuchen Wu. 2021. Scalable Bottom-up Hierarchical Clustering. *Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery Data Mining*.
- Nicholas Monath, A. Kobren, A. Krishnamurthy, Michael R. Glass, and A. McCallum. 2019. Scalable Hierarchical Clustering with Tree Grafting. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery Data Mining*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- 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.
- Vered Shwartz, Gabriel Stanovsky, and Ido Dagan. 2017. [Acquiring predicate paraphrases from news tweets](#). In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pages 155–160, Vancouver, Canada. Association for Computational Linguistics.
- 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.
- Wei Wu, Fei Wang, Arianna Yuan, Fei Wu, and Jiwei Li. 2020. [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.
- Bishan Yang, Claire Cardie, and Peter Frazier. 2015. [A hierarchical distance-dependent Bayesian model for event coreference resolution](#). *Transactions of the Association for Computational Linguistics*, 3:517–528.
- Yutao Zeng, Xiaolong Jin, Saiping Guan, Jiafeng Guo, and Xueqi Cheng. 2020. [Event coreference resolution with their paraphrases and argument-aware embeddings](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3084–3094, Barcelona, Spain (Online). International Committee on Computational Linguistics.

A Appendix

A.1 Implementation Details

Our model includes 14M parameters and is implemented in PyTorch (Paszke et al., 2019), using HuggingFace’s library (Wolf et al., 2020) and the Adam optimizer (Kingma and Ba, 2014). The layers of the models are initialized with Xavier Glorot method (Glorot and Bengio, 2010). We manually tuned the standard hyperparameters, presented in Table 5 on the event coreference task and keep them unchanged for entity and ALL settings. Table 6 shows specific parameters, such as the maximum span width, the pruning coefficient λ and the stop criterion τ for the agglomerative clustering, that we tuned separately for each setting to maximize the CoNLL F1 score on its corresponding development set.

Hyperparameter	Value
Batch size	32
Dropout	0.3
Learning rate	0.0001
Optimizer	Adam
Hidden layer	1024

Table 5: Shared hyperparameters across the different models.

	Max span width	λ	τ
Event	10	0.25	0.75
Entity	15	0.35	0.75
ALL	15	0.4	0.75

Table 6: Specific hyperparameters for each mention type; λ is the pruning coefficient and τ is the threshold for the agglomerative clustering.

A.2 Dataset

ECB+⁴ is an extended version of the EventCoref-Bank (ECB) (Bejan and Harabagiu, 2010) and EECB (Lee et al., 2012), whose statistics are shown in Table 7. The dataset is composed of 43 *topics*, where each topic corresponds to a famous news event (e.g Someone checked into rehab). In order to introduce some complexity and to limit the use of lexical features, each topic is constituted by a collection of texts describing two different event

instances of the same event type, called *subtopic*. For example, the first topic corresponding to the event “Someone checked into rehab” is composed of event mention of the event “Tara Reid checked into rehab” and “Lindsay Lohan checked into rehab” which are obviously annotated into different coreference cluster. Documents in ECB+ are in English. Since ECB+ is an *event-centric* dataset, entities are annotated only if they participate in events. In this dataset, event and entity coreference clusters are denoted separately.

	Train	Validation	Test
# Topics	25	8	10
# Documents	594	196	206
# Mentions	3808/4758	1245/1476	1780/2055
# Singletons	1116/814	280/205	632/412
# Clusters	1527/1286	409/330	805/608

Table 7: ECB+ statistics. The slash numbers for # Mentions, # Singletons and # Clusters represent event/entity statistics. As recommended by the authors in the release note, we follow the split of Cybulska and Vossen (2015) that use a curated subset of the dataset.

⁴<http://www.newsreader-project.eu/results/data/the-ecb-corpus/>