

A Controlled Reevaluation of Coreference Resolution Models

Ian Porada*, Xiyuan Zou*, Jackie Chi Kit Cheung

McGill University, Mila - Quebec AI Institute

{ian.porada, xiyuan.zou}@mail.mcgill.ca, jackie.cheung@mcgill.ca

Abstract

All state-of-the-art coreference resolution (CR) models involve finetuning a pretrained language model. Whether the superior performance of one CR model over another is due to the choice of language model or other factors, such as the task-specific architecture, is difficult or impossible to determine due to lack of a standardized experimental setup. To resolve this ambiguity, we systematically evaluate five CR models and control for certain design decisions including the pretrained language model used by each. When controlling for language model size, encoder-based CR models outperform more recent decoder-based models in terms of both accuracy and inference speed. Surprisingly, among encoder-based CR models, more recent models are not always more accurate, and the oldest CR model that we test generalizes the best to out-of-domain textual genres. We conclude that controlling for the choice of language model reduces most, but not all, of the increase in F1 score reported in the past five years.

Keywords: coreference resolution, reproducibility, confounding factors

1. Introduction

Coreference resolution (CR) is the task of clustering spans of text in a discourse that refer to the same entity or event (Hobbs, 1978, 1979).

All state-of-the-art CR models incorporate supervised finetuning of a pretrained language model. See Poesio et al. (2023) for a comprehensive survey. We study the question:

To what extent are recent improvements in CR due to use of a more powerful language model as opposed to other design decisions?

To answer this question, we evaluate existing CR models, controlling for the language model used in each. We find that at comparable language model sizes, encoder-based CR models outperform decoder-based models in terms of accuracy, inference speed, and memory usage.

This finding holds when we scale encoder-based models to sizes larger than existing work. For instance, when scaled to 1.5B parameters the encoder-based LingMess model (Otmazgin et al., 2023) achieves 82.5 CoNLL F1 score on the CoNLL-2012 Shared Task test set. This is the same score as the decoder-based ASP model at 11B parameters (Liu et al., 2022).

When we control for additional factors, such as the search space of supervised training epochs, we find that the oldest model tested, C2F (Lee et al., 2018), performs competitively with all other encoder-based models. Furthermore, C2F generalizes better to out-of-domain textual genres than all models of comparable size.

Based on our results, we conclude that controlling for the choice of language model reduces

most, but not all, of the increase in F1 score reported in the past five years. Many improvements in CR model accuracy in may therefore be attributable to use of a stronger language model as opposed to other design decisions. This finding suggests that future proposals intended to improve CR accuracy should carefully consider possible causes in order to better understand to what extent improvements are attributable to proposed architectural changes.

The five CR models that we test in our experiments are: 1) C2F; 2) S2E (Kirstain et al., 2021); 3) WLC (Dobrovolskii, 2021); 4) LingMess; and, 5) Link-Append (Bohnet et al., 2023). See §3.1 for model details. We focus our evaluation on English-language, document-level, nominal-entity coreference resolution as formulated in OntoNotes (Hovy et al., 2006). We additionally test generalization to the GAP (Webster et al., 2018) and OntoGUM (Zhu et al., 2021b) corpora. See §3.3 for dataset details.

The main contributions of this work are: **1)** We reimplement five state-of-the-art CR models, reproducing the original, published results. **2)** We show that more recent CR models are not always more accurate when we control for the choice of language model and the hyperparameter search space. **3)** We scale up encoder-based CR models to 1.5B parameters and find that they outperform decoder-based models of comparable size.

2. Related Work

Existing work has presented controlled comparisons of published models, notably in the case of language model pretraining (Melis et al., 2018; Nityasya et al., 2023). While similar in motivation to our work, we focus specifically on CR, with em-

*Equal contribution.

phasis on the choice of language model.

Previous work has attempted to better understand state-of-the-art CR models based on targeted error analyses (Stoyanov et al., 2009; Lu and Ng, 2020) and model generalization across datasets (Toshniwal et al., 2021; Žabokrtský et al., 2022; Porada et al., 2023); however, limited work has systematically evaluated existing model design. Martschat and Strube (2015) performed a comprehensive evaluation of CR models, evaluating pre-neural architectures. More recent work has focused on analyzing particular steps in CR model architecture design (Toshniwal et al., 2020; Xu and Choi, 2020; Lai et al., 2022).

Many existing CR models include ablations allowing for pairwise comparison between certain models (Dobrovolskii, 2021; Liu et al., 2022); however, our study provides novel insights in that we train all models with the same competitive language model, where possible. Our results reveal the accuracy of C2F to be underestimated in existing comparisons.

3. Methods

In this section, we introduce details of each model, the factors that we control for, and the datasets used for evaluation.

3.1. Models

We focus our evaluation on four encoder-based models and one decoder-based model, all of which reported state-of-the-art accuracy at their respective time of publication.

3.1.1. Encoder-based Models

C2F The C2F (Coarse-to-Fine) model (Lee et al., 2018) is an extension of the earlier E2E (Lee et al., 2017). E2E functions by encoding spans of text as contextualized vectors, and then pairwise scoring these vector representations using a task-specific head. C2F extends E2E with an intermediate bilinear scoring function that filters down the number of spans considered. Lee et al. (2018) also proposed higher-order inference (HOI) and originally evaluated both E2E+C2F and E2E+C2F+HOI. We specifically consider only E2E+C2F, which we refer to as the C2F model, because HOI is known to marginally affect performance (Xu and Choi, 2020). While the original C2F proposal did not finetune the language model encoder, we follow the hyperparameters of Joshi et al. (2019), which was the first work to do so.

S2E The S2E (Start-to-End) model (Kirstain et al., 2021) is based on C2F with the main distinction that span representations are created using only the embeddings of the first and last token in the span. By contrast, span representations

Model	Language Model	Prior	Ours
C2F	BERT-large	76.9	77.1
S2E	Longformer-large	80.3	80.4
WLC	RoBERTa-large	81.0	81.0
LingMess	Longformer-large	81.4	81.4
Link-Append	mT5-XXL	83.3	83.3

Table 1: Test set performance on OntoNotes (CoNLL-2012) as reported in prior work (“prior”) and for our reimplementation of the model (“ours”). For each model, we use the best configuration of the respective prior work.

in C2F are a weighted sum of all token embeddings. S2E was motivated as requiring less memory while maintaining accuracy.

WLC The WLC model (Word-Level Coreference; Dobrovolskii, 2021) is also based on C2F. In WLC, individual tokens are scored as candidate mentions, each representing the headword of some span. WLC was motivated as being faster than C2F while also more accurate when using a RoBERTa language model (Liu et al., 2019).

LingMess Otmazgin et al. (2023) proposed LingMess as a direct extension of S2E by increasing the number of task-specific, span-scoring heads. LingMess significantly improved accuracy when using a Longformer language model (Beltagy et al., 2020) as the encoder.

3.1.2. Decoder-based Models

Two competitive decoder-based CR models have been proposed in the literature: ASP (Liu et al., 2022) and Link-Append. Of these we evaluate Link-Append, which reported a higher F1 score on OntoNotes, although we compare against ASP using published numbers where possible.

Link-Append Bohnet et al. (2023) proposed Link-Append, a method to finetune a language model for CR purely as a sequence generation task. For each sentence in the input document, a language model is trained to output a string of all coreferring mentions in the sentence. When used to finetune the 13B parameter mT5 language model (Xue et al., 2021), Link-Append demonstrated state-of-the-art accuracy.

3.1.3. Implementation

We reimplement all models using Hugging Face’s Transformers library (Wolf et al., 2020) following each model’s original hyperparameters.* This standardized setting allows for comparison of empirical inference speed and memory usage. Additionally, we can easily control for factors such

*Our code is available at <https://github.com/ianporada/coref-reeval>

as the language model used, whereas existing implementations often require manually-run preprocessing steps. We base our reimplementation on the original model code when available.

We verify our reimplementation of each model by comparing against the model accuracy for the best published configurations (Table 1). As we do not have the resources to train the Link-Append model at the original 13 billion parameter size, we verify our processing steps by running inference on the released model weights. We also clarify preprocessing details with the original authors.

3.2. Considered Factors

Our evaluation tests whether improved accuracy of a CR model is due to the proposed model design versus other factors. The main factor that we consider is the language model used in the initialization of each CR model. More recent models have also made changes to hyperparameters, such as the number of supervised training epochs, which may be responsible for improvements in performance. To control for these factors, we evaluate all models when tuned over the same hyperparameter space, a superset of existing search spaces.

3.2.1. Language Model

Encoder For encoder-based CR models, we train all models using the same competitive language model encoder, DeBERTa (He et al., 2021), at base and large sizes. We choose DeBERTa due to its competitive performance on the SuperGLUE benchmark (Wang et al., 2019).

Encoder v.s. Decoder We do not directly compare encoder versus decoder-based CR models using the exact same language model because these classes of CR models in turn rely on language models of different architectures. Instead, we control for language model size. We train the best performing encoder-based model, LingMess, using DeBERTa encoders ranging from 138M to 1.5B parameters. We then compare against Link-Append and ASP decoder-based CR models trained with language models ranging from 300M to 13B parameters.

3.2.2. Hyperparameter Search Space

Finetuning Epochs More recent models have been trained for longer: while Joshi et al. (2019) trained C2F for 20 epochs, LingMess has been trained for 129 epochs, and the 100k steps used to train Link-Append amounts to around 160 epochs. To control for this change as a cause of improved accuracy, we search over training epochs in {25, 50, 125} for each model.

Task-specific Head Size Post-C2F, encoder-based CR models use a larger hidden-layer size for task-specific heads (`ffnn_size`). To control for

DeBERTa-base (138M)					
Method	ON	OG	GAP	Mem.	Time
C2F	77.3	63.5	77.4	6.2	71.6
S2E	78.6	63.8	78.9	1.3	20.9
WLC	79.1	64.1	78.7	1.4	25.9
LingMess	79.5	64.3	79.4	1.7	54.7
DeBERTa-large (405M)					
Method	ON	OG	GAP	Mem.	Time
C2F	80.8	66.9	79.1	8.7	129.3
S2E	80.5	65.9	79.9	2.7	45.3
WLC	81.0	66.2	79.1	2.8	47.4
LingMess	81.7	66.4	79.4	3.1	81.6

Table 2: Dev. set performance using a DeBERTa encoder of the given size. ON and OG are evaluated by CoNLL F1 score, GAP by F1 score. Mem. is max memory at batch size of one (GB), and time is inference speed in ms/doc at max batch size, both calculated w.r.t. inference on ON using a single 80GB A100 GPU.

the possible impact of this increase, we search over `ffnn_size` \in {1024, 2048, 3072, 4096} for every encoder-based model. This search space includes all sizes used by existing models.

Input Size At training time, C2F reduces the length of documents to a small, fixed number of sentences in order to reduce memory usage, whereas other models do not have this constraint. We remove this constraint by training C2F on the maximum document length that fits into memory.

3.3. Datasets

OntoNotes (ON) As in existing work, we train and evaluate all models on the coreference annotations in the English-language portion of the OntoNotes 5.0 dataset (Weischedel et al., 2013). We specifically use the CoNLL-2012 Shared Task v4 dataset splits (Pradhan et al., 2012). For evaluation, we use the official CoNLL-2012 scorer. The train/validation/test splits are 1940/343/348 document parts, respectively.

OntoGUM (OG) We additionally consider generalization to out-of-domain genres by evaluating models on the OntoGUM (OG) dataset (Zhu et al., 2021b). OG is composed of the coreference annotations in the English-language GUM corpus (Zeldes, 2017) transformed heuristically to follow OntoNotes annotation guidelines (Zhu et al., 2021a). When evaluating on OG, we always evaluate on all 213 documents in the dataset.

GAP Finally, we evaluate models on the 2000-example validation set of the GAP dataset (Webster et al., 2018) as a targeted evaluation of pronominal anaphora. GAP consists of pronouns

Model	LM	Size	ON	OG	GAP	Mem.	Time
C2F		452M	81.2 [†]	66.9	78.1	8.7	129.3
S2E	DeBERTa _L	465M	80.5	65.9 [†]	79.9	2.7	45.3
WLC		664M	81.1 [†]	65.8 [†]	78.8	2.8	47.4
LingMess		561M	81.7	66.4	79.4	3.1	81.6
Link-Append	mT5 _B	582M	66.2	45.8	63.3	5.9	5.2e4
	mT5 _L	1.23B	66.5	45.5	64.9	12.3	1.4e5

Table 3: Dev. set accuracy when controlling for all factors. ON and OG are evaluated by CoNLL F1 score, GAP by F1 score. Mem. is the max memory usage for inference at batch size one (GB), and time is inference speed (ms/doc) at max batch size, both w.r.t. ON inference using a single 80 GB A100 GPU. F1 scores without † are statistically significantly different from all other scores based on a permutation test with $\alpha = 0.1$ and 10k permutations following Chinchor (1995).

in English Wikipedia annotated for coreference with respect to two preceding noun phrases. We score models using the official scorer.

4. Experiments

We now present results controlling for language model size and hyperparameter search space. We find that controlling for these factors considerably narrows the gap in accuracy between models.

More Recent Encoder-based Models Are Not Always More Accurate First, we test all CR models in their original configuration, with the minimal change that we use a DeBERTa encoder (Table 2). When using DeBERTa-large, the difference in CoNLL F1 on OntoNotes between all models reduces to less than 1 point absolute, and older models (C2F and S2E) are most accurate out-of-domain. Every encoder-based model also performs at or better than its best reported configuration in the literature.

We also control for the hyperparameter search space by comparing all models in their best hyperparameter configuration in Table 3. In this case, the difference in performance between the oldest C2F and most recent LingMess encoder models reduces to 0.5 points absolute on OntoNotes. For all models in Table 3, full precision, recall, and test set results are available in our project’s GitHub repository, as are human-readable model predictions in CoNLL format.

Encoder-based Models Are More Accurate than Decoder-based Models at Comparable Sizes Table 3 shows the accuracy and inference time results for all encoder-based models versus Link-Append trained at approximately the same size. At this size, encoder-based models are substantially both more accurate and faster than Link-Append.

Accuracy of models trained over a range of sizes are shown in Figure 1. Up to the largest LingMess model trained, 1.5B parameters, LingMess outperforms decoder-based methods.

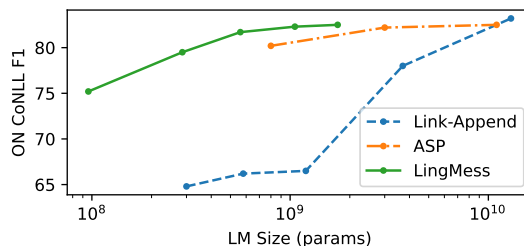


Figure 1: OntoNotes (ON) test set CoNLL F1 score for CR models trained with language models at multiple scales. Link-Append is trained with mT5, ASP with Flan-T5, and LingMess with DeBERTa. ASP scores, as well as the two largest Link-Append model scores, were reported by the respective authors. Bohnet (2023) also noted that Link-Append has an unexpected, apparently “emergent”, scaling.

4.1. Discussion

Our results may be surprising because CR models more recent than C2F, such as WLC and S2E, are often reported to be more accurate than C2F when comparisons are presented (Chai and Strube, 2022; Liu et al., 2022; Bohnet et al., 2023); yet, we found the opposite to be true when controlling for the choice of language model. We suggest that future work explicitly considers these factors when presenting comparisons of CR architectures.

In some cases, less accurate models are substantially more memory or time efficient (Table 3). While these dimensions have been considered in previous work (Kirstain et al., 2021), even these efficient models have been motivated by their raw accuracy. To this end, we suggest more holistic evaluations of coreference beyond a single accuracy number, especially considering the fact that high OntoNotes F1 does not necessarily translate to good out-of-domain performance.

In our controlled comparison, the WLC model is never Pareto-optimal based on accuracy, run-

Former President Clinton drew cheers from admirers at a US Air Force base outside Washington **today** before departing for
gold_0 pred_0 gold_1 pred_1

his home in New York , and although **Mr. Clinton** is out of **office** , **he** says **he** 'll still be around . So you see that sign there says
gold_0 pred_0 gold_0 pred_0 pred_2 gold_0 pred_0 gold_0 pred_0

, `` Please do n't go '' . **I** left the White House , but **I** 'm still here . Go anywhere . Prior to leaving **office today Mr. Clinton**
gold_0 pred_0 gold_0 pred_0 pred_2 gold_1 pred_1 gold_0 pred_0

issued more than one hundred presidential pardons . Among those pardoned was Susan Mcdoogle , a former business partner of
Mr. Clinton who became a key figure in the White Water Real Estate investigation that plagued the **Clinton** presidency .
gold_0 pred_0 gold_0 pred_0

Figure 2: An example model output on the OntoNotes development set. “Gold” spans denote the ground-truth annotations, and “pred” spans denote predictions by the Link-Append mT5-XXL model.

time, and memory usage. All other encoder-based models are optimal in at least one of these three dimensions.

5. Conclusion

We reevaluate existing CR models, controlling for factors such as the choice of language model and number of training epochs. Among encoder-based models, the relative performance gap on OntoNotes narrows to 0.5 absolute CoNLL F1 score, and the oldest model tested (C2F) generalizes the best to out-of-domain textual genres. At all scales tested—up to 1.5B parameters—encoder-based CR models outperform decoder-based models of comparable model size; notably, LingMess trained with a 1.5B DeBERTa encoder has the same accuracy as decoder-based models over seven times larger in size.

Acknowledgements

The authors acknowledge the material support of NVIDIA in the form of computational resources. Ian Porada is supported by a fellowship from the Fonds de recherche du Québec (FRQ). Xiyuan Zou is supported by a McGill Science Undergraduate Research Award (SURA). Jackie Chi Kit Cheung is supported by the Canada CIFAR AI Chair program. We thank Bernd Bohnet for help reimplementing the Link-Append model.

6. Bibliographical References

Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv:2004.05150*.

Bernd Bohnet. 2023. [Multilingual coreference resolution with innovative seq2seq models](#). Invited Talk at The Sixth Workshop on Computational Models of Reference, Anaphora and Coreference (CRAC 2023).

Bernd Bohnet, Chris Alberti, and Michael Collins. 2023. [Coreference resolution through a seq2seq transition-based system](#). *Transactions of the Association for Computational Linguistics*, 11:212–226.

Haixia Chai and Michael Strube. 2022. [Incorporating centering theory into neural coreference resolution](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2996–3002, Seattle, United States. Association for Computational Linguistics.

Nancy Chinchor. 1995. [Statistical significance of MUC-6 results](#). In *Sixth Message Understanding Conference (MUC-6): Proceedings of a Conference Held in Columbia, Maryland, November 6-8, 1995*.

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.

Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.

- Jerry R Hobbs. 1978. Resolving pronoun references. *Lingua*, 44(4):311–338.
- Jerry R Hobbs. 1979. Coherence and coreference. *Cognitive science*, 3(1):67–90.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. [OntoNotes: The 90% solution](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 57–60, New York City, USA. Association for Computational Linguistics.
- Mandar Joshi, Omer Levy, Luke Zettlemoyer, and Daniel Weld. 2019. [BERT for coreference resolution: Baselines and analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5803–5808, Hong Kong, China. 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.
- Tuan Manh Lai, Trung Bui, and Doo Soon Kim. 2022. End-to-end neural coreference resolution revisited: A simple yet effective baseline. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8147–8151. IEEE.
- 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.
- Tianyu Liu, Yuchen Eleanor Jiang, Nicholas Monath, Ryan Cotterell, and Mrinmaya Sachan. 2022. [Autoregressive structured prediction with language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 993–1005, Abu Dhabi, United Arab Emirates. 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 pre-training approach](#). *arXiv:1907.11692*.
- Jing Lu and Vincent Ng. 2020. [Conundrums in entity coreference resolution: Making sense of the state of the art](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6620–6631, Online. Association for Computational Linguistics.
- Sebastian Martschat and Michael Strube. 2015. [Latent structures for coreference resolution](#). *Transactions of the Association for Computational Linguistics*, 3:405–418.
- Gábor Melis, Chris Dyer, and Phil Blunsom. 2018. [On the state of the art of evaluation in neural language models](#). In *International Conference on Learning Representations*.
- Made Nindyatama Nityasya, Haryo Wibowo, Alham Fikri Aji, Genta Winata, Radityo Eko Prasoj, Phil Blunsom, and Adhiguna Kuncoro. 2023. [On “scientific debt” in NLP: A case for more rigour in language model pre-training research](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8554–8572, Toronto, Canada. Association for Computational Linguistics.
- Shon Otmazgin, Arie Cattan, and Yoav Goldberg. 2023. [LingMess: Linguistically informed multi expert scorers for coreference resolution](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2752–2760, Dubrovnik, Croatia. Association for Computational Linguistics.
- Massimo Poesio, Juntao Yu, Silviu Paun, Abdulrahman Aloraini, Pengcheng Lu, Janosch Haber, and Derya Cokal. 2023. Computational models of anaphora. *Annual Review of Linguistics*, 9:561–587.
- Ian Porada, Alexandra Olteanu, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2023. [Investigating failures to generalize for coreference resolution models](#). *arXiv:2303.09092*.

- 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.
- Veselin Stoyanov, Nathan Gilbert, Claire Cardie, and Ellen Riloff. 2009. [Conundrums in noun phrase coreference resolution: Making sense of the state-of-the-art](#). In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 656–664, Suntec, Singapore. Association for Computational Linguistics.
- Shubham Toshniwal, Haoyue Shi, Bowen Shi, Lingyu Gao, Karen Livescu, and Kevin Gimpel. 2020. [A cross-task analysis of text span representations](#). In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 166–176, Online. Association for Computational Linguistics.
- Shubham Toshniwal, Patrick Xia, Sam Wiseman, Karen Livescu, and Kevin Gimpel. 2021. [On generalization in coreference resolution](#). In *Proceedings of the Fourth Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 111–120, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. [Super-glue: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- 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.
- Liyan 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.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Zdeněk Žabokrtský, Miloslav Konopík, Anna Nedoluzhko, Michal Novák, Maciej Ogrodniczuk, Martin Popel, Ondřej Pražák, Jakub Sido, Daniel Zeman, and Yilun Zhu. 2022. [Findings of the shared task on multilingual coreference resolution](#). In *Proceedings of the CRAC 2022 Shared Task on Multilingual Coreference Resolution*, pages 1–17, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Amir Zeldes. 2017. [The GUM corpus: Creating multilayer resources in the classroom](#). *Language Resources and Evaluation*, 51(3):581–612.
- Yilun Zhu, Sameer Pradhan, and Amir Zeldes. 2021a. [Anatomy of OntoGUM—Adapting GUM to the OntoNotes scheme to evaluate robustness of SOTA coreference algorithms](#). In *Proceedings of the Fourth Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 141–149, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yilun Zhu, Sameer Pradhan, and Amir Zeldes. 2021b. [OntoGUM: Evaluating contextualized SOTA coreference resolution on 12 more genres](#). 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 461–467, Online. Association for Computational Linguistics.

7. Language Resource References

Weischedel, Ralph and others. 2013. *OntoNotes Release 5.0*. Defense Advanced Research Projects Agency, GALE Program. Linguistic Data Consortium, LDC2013T19, ISLRN 151-738-649-048-2. [\[link\]](#).