

LINGMESS: Linguistically Informed Multi Expert Scorers for Coreference Resolution

Shon Otmazgin¹ Arie Cattan¹ Yoav Goldberg^{1,2}

¹Computer Science Department, Bar Ilan University

²Allen Institute for Artificial Intelligence

{shon711, arie.cattan, yoav.goldberg}@gmail.com

Abstract

Current state-of-the-art coreference systems are based on a single pairwise scoring component, which assigns to each pair of mention spans a score reflecting their tendency to corefer to each other. We observe that different kinds of mention pairs require different information sources to assess their score. We present LINGMESS, a linguistically motivated categorization of mention-pairs into 6 types of coreference decisions and learn a dedicated trainable scoring function for each category. This significantly improves the accuracy of the pairwise scorer as well as of the overall coreference performance on the English Ontonotes coreference corpus and 5 additional datasets.¹

1 Introduction

Coreference resolution is the task of clustering textual mentions that refer to the same discourse entity. This fundamental task requires many decisions. In this work, we argue that different *kinds* of decisions involve different challenges. To illustrate that, consider the following text:

“Lionel Messi has won a record seven Ballon d’Or awards. He signed for Paris Saint-Germain in August 2021. “I would like to thank my family”, said the Argentinian footballer. Messi holds the records for most goals in La Liga”

To correctly identify that the pronoun “He” refers to “Lionel Messi”, models need to model the discourse, while linking “my” to “I” may rely more heavily on morphological agreement. Likewise, linking “the Argentinian footballer” to “Lionel Messi” requires world knowledge, while linking “Messi” to “Lionel Messi” may be achieved by simple lexical heuristics.

¹The codebase to train and run LINGMESS is available in <https://github.com/shon-otmazgin/lingmess-coref>. Also, our recent F-COREF Python package (Otmazgin et al., 2022) includes a simple and efficient implementation of LINGMESS in <https://github.com/shon-otmazgin/fastcoref>.




Figure 1: Architecture of our multi expert model. Given two spans “Lionel Messi” and “He”, we sum four scores: individual mention scores (black), $f_m(\text{“Lionel Messi”})$, $f_m(\text{“He”})$, and pairwise scores, shared antecedent score (white) $f_a(\text{“Lionel Messi”, “He”})$ and the relevant “expert” score (blue) $f_a^{\text{PRON-ENT}}(\text{“Lionel Messi”, “He”})$.

Indeed, pre-neural coreference resolution works often considered the types of a mention-pair, either by incorporating this information as model features, or by tailoring specific rules or specific models for each mention pair (see related work section).

However, neural-network based coreference models are all based on a *single* pairwise scorer that is shared for all mention pairs, regardless of the different challenges that needs to be addressed by each pair type (Lee et al., 2017, 2018; Joshi et al., 2019; Kantor and Globerson, 2019; Joshi et al., 2020; Xu and Choi, 2020; Xia et al., 2020; Toshniwal et al., 2020; Thirukovalluru et al., 2021; KIRSTAIN et al., 2021; Dobrovolskii, 2021).

In this work, we suggest that modeling different mention pairs by different sub-models (in our case, different learned scoring functions) depending on their types is beneficial also for neural models. We identify a set of decisions: (a) linking compatible

Category	Co-referring example	Non Co-referring example
PRON-PRON-C	<i>A couple of my law clerks were going to ... and I was afraid I was going to...</i>	<i>The Lord God said to my Lord: “Sit by me at my right side , and I will put your enemies ... ”</i>
PRON-PRON-NC	<i>“I made a similar line and I produced it cheaper” , he says.</i>	<i>She is my Goddess ...</i>
ENT-PRON	<i>Spain, Argentina, Thailand and Indonesia were doing too little to prevent ... across their borders.</i>	<i>Tonight, to kick off the effort, CNN will premiere its first prime - time newscast in years.</i>
MATCH	<i>... says Paul Amos, CNN executive vice president for programming. Accordingly, CNN is ...</i>	<i>Hertz and Avis can not benefit Budget’s programs, ” said Bob Wilson, Budget’s vice president ...</i>
CONTAINS	<i>He reportedly showed DeLay a videotape that made him weep. Tom DeLay then ...</i>	<i>Give SEC authority to halt securities trading, (also opposed by new SEC chairman) ...</i>
OTHER	<i>They also saw the two men who were standing with him. When Moses and Elijah were leaving ...</i>	<i>The company is already working on its own programming ... the newspaper said.</i>

Table 1: Example of each category, taken from Ontonotes development set. We define the categories of mention pairs as follows. PRON-PRON-C: compatible pronouns based on their attributes such as gender, number and animacy (see Appendix C for more details), PRON-PRON-NC: incompatible pronouns, ENT-PRON: a pronoun and another span, MATCH: non-pronoun spans with the same content words, CONTAINS: one contains the content words of the other, OTHER: all other pairs. Content words exclude stop words, see Appendix C for the full list of stop words.

pronouns (PRON-PRON-C); (b) linking incompatible pronouns (PRON-PRON-NC); (c) linking pronouns to entities (ENT-PRON); (d) linking entities which share the exact lexical form (MATCH); (e) linking entities where the lexical form of one contains the lexical form of the other (CONTAINS); (f) other cases (OTHER). Each of these classes is easy to identify deterministically, each contains both positive and negative instances, and each could benefit from a somewhat different decision process. Table 1 demonstrates the classes.²

We present **Linguistically Informed Multi Expert Scorers (LINGMESS)**, a coreference model which categorizes each pairwise decision into one of these classes, and learns, in addition to a shared scoring function, also a separate scoring function for each pair type. At inference time, for each pair of mentions being scored, we deterministically identify the pair’s type, and use the corresponding scoring function.³

Specifically, we extend the recent *s2e*’s model (Kirstain et al., 2021) by adding per-category scoring, but the method is general and may work with other coreference models as well. As illustrated in Figure 1, the final coreference score between two spans is composed—in addition to the individual mention scores—of two pairwise scores: a shared antecedent-compatibility score and an “ex-

pert” antecedent compatibility score which depends on the linguistic-type of the pair.

We show that this significantly improves the coreference performance on Ontonotes (Pradhan et al., 2012) and 5 additional datasets. We also inspect the performance of the model for each category separately, showing that some classes improve more than others. This analysis provides a finer-grained understanding of the models and points out directions for future research.

2 Background: the *s2e* Model

The *s2e* model (Kirstain et al., 2021) achieves the current best coreference scores among all practical neural models.⁴

Given a sequence of tokens x_1, \dots, x_n , each mention pair (c, q) is scored using a scoring function $F_s(c, q)$ ⁵ described below, where c is a “candidate span” and q is a “query span” which appears before c in the sentence. The span encodings are based on contextualized word embeddings obtained by a Longformer encoder, see Kirstain et al. (2021) for details. These pairwise scores are then used to form coreference chains (see “inference” below).

²More fine-grained distinctions are of course also possible, but we leave exploration of them to future work.

³For computational efficiency on the GPU, we find it beneficial to compute all the scoring functions and mask away the not needed values.

⁴We define “practical models” as those that require a constant number transformer-based document encodings per passage, as opposed to a constant number of document encodings per mention. The CorefQA model (Wu et al., 2020) achieves a substantially higher score, but requires to run a separate BERT inference for each mention, making it highly impractical.

⁵S state for SHARED

The scoring function F_s is further decomposed:

$$F_s(c, q) = \begin{cases} f_m(c) + f_m(q) + f_a(c, q) & c \neq \varepsilon \\ 0 & c = \varepsilon \end{cases}$$

where ε is the null antecedent, and f_m and f_a are parameterized functions, scoring each individual span (f_m) and the pairwise interaction (f_a).

For each possible mention q , the learning objective optimizes the sum of probabilities over the true antecedent \hat{c} of q :

$$L_s(q) = \log \sum_{\hat{c} \in \mathcal{C}(q) \cap \text{GOLD}(q)} P_s(\hat{c} | q)$$

where $\mathcal{C}(q)$ is the set of all candidate antecedents⁶ with a null antecedent ε . $\text{GOLD}(q)$ is the set of the true antecedents of q . $P_s(c | q)$ is computed as a softmax over $F_s(c, q)$ scores for c values in $\mathcal{C}(q)$:

$$P_s(\hat{c} | q) = \frac{\exp F_s(\hat{c}, q)}{\sum_{c' \in \mathcal{C}(q)} \exp F_s(c', q)}$$

3 LINGMESS

Clustering coreferring entities typically involves many different phenomena, which we argue should be addressed in a different manner. Therefore, our core contribution is proposing to allocate a dedicated scorer $f_a^t(c, q)$ for each phenomenon type t , in addition to the shared pairwise scorer $f_a(c, q)$. The overall architecture of our model is shown in Figure 1.

Concretely, we extend the *s2e* model with six additional antecedent scorers f_a^t where $t \in \{\text{PRON-PRON-C}, \text{PRON-PRON-NC}, \text{ENT-PRON}, \text{MATCH}, \text{CONTAINS}, \text{OTHER}\}$, the six categories we list in Table 1.

The pairwise scoring function now becomes:

$$F(c, q) = \begin{cases} f_m(c) + f_m(q) + f(c, q) & c \neq \varepsilon \\ 0 & c = \varepsilon \end{cases}$$

$$f(c, q) = f_a(c, q) + f_a^{T(c, q)}(c, q)$$

where $T(c, q)$ is a deterministic, rule-based function to determine the category t of the pair (c, q) . The pairwise scoring function $f(c, q)$ scoring c as the antecedent of q , is now composed of a shared scorer $f_a(c, q)$ and an ‘‘expert’’ scorer $f_a^t(c, q)$ which differs based on the type of the pair c, q . Each of the seven pairwise scoring functions (f_a

⁶All spans before q that passed some pruning threshold.

and the six f_a^t) is parameterized separately using its own set of matrices. The transformer-based encoder and the mention scorer f_m are shared between all the antecedent scorers. See Appendix A.2 for full model architecture.

Training For each span q , our model optimizes the objective function L_{COREF} over the sum of probabilities of all true antecedents of q :

$$L_{\text{COREF}}(q) = \log \sum_{\hat{c} \in \mathcal{C}(q) \cap \text{GOLD}(q)} P(\hat{c} | q)$$

Here, $P(\hat{c} | q)$ is a softmax over $F(\hat{c}, q)$ scores, that is our new score function described in Figure 1.

$$P(\hat{c} | q) = \frac{\exp F(\hat{c}, q)}{\sum_{c' \in \mathcal{C}(q)} \exp F(c', q)}$$

This scorer is also the one used in inference. However, this objective does not explicitly push each category (‘‘expert’’) to specialize. For example, for the PRON-PRON-C cases, it would be useful to explicitly train the model to distinguish between the possible antecedents of that type only (without regarding other antecedents), as well as to explicitly distinguish between a pronoun antecedent and a null antecedent. To this end, we extend the training objective by also training each ‘‘expert’’ separately:

$$L_t(q) = \log \sum_{\hat{c} \in \mathcal{C}_t(q) \cap \text{GOLD}(q)} P_t(\hat{c} | q)$$

$$P_t(\hat{c} | q) = \frac{\exp F_t(\hat{c}, q)}{\sum_{c' \in \mathcal{C}_t(q)} \exp F_t(c', q)}$$

$$F_t(c, q) = \begin{cases} f_m(c) + f_m(q) + f_a^t(c, q) & c \neq \varepsilon \\ 0 & c = \varepsilon \end{cases}$$

Note that for $L_t(q)$ we replace $\mathcal{C}(q)$ with $\mathcal{C}_t(q)$, considering only the potential antecedents that are compatible with the span q for the given type. For example, for L_{MATCH} and a span q , we will only consider candidates c which appear before q with the exact same content words as q . Our final objective for each mention span q is thus:

$$L(q) = L_{\text{COREF}}(q) + L_{\text{EXPERTS}}(q)$$

$$L_{\text{EXPERTS}}(q) = \sum_t L_t(q) + L_s(q)$$

	MUC			B ³			CEAF _{φ₄}			LEA			Avg. F1
	R	P	F1	R	P	F1	R	P	F1	R	P	F1	
<i>s2e</i>	85.2	86.6	85.9	77.9	80.3	79.1	75.4	76.8	76.1	75.8	78.3	77.0	80.3
LINGMESS	85.1	88.1	86.6	78.3	82.7	80.5	76.1	78.5	77.3	76.3	80.9	78.5	81.4

Table 2: Performance on the test set of the English OntoNotes 5.0 dataset. The averaged F1 of MUC, B³, CEAF_φ is the main evaluation metric. Our model outperforms the *s2e* model (Kirstain et al., 2021) by 1.1 CoNLL F1. The performance gain is statistically significant according to a non-parametric permutation test ($p < 0.05$).

	<i>s2e</i>	LINGMESS
WikiCoref (Ghaddar and Langlais, 2016)	59.7	62.6
GAP (Webster et al., 2018)	88.3	89.6
WinoGender (Rudinger et al., 2018)	70.5	77.3
WinoBias (Zhao et al., 2018)	84.3	85.1
BUG (Levy et al., 2021)	72.2	74.6

Table 3: Performance on the test set of various coreference datasets. The reported metrics are CoNLL F1 for WikiCoref, F1 for GAP and Accuracy for WinoGender, WinoBias and BUG.

	<i>s2e</i>			LINGMESS		
	P	R	F1	P	R	F1
PRON-PRON-C	88.8	71.3	79.1	88.0	85.1	86.5
PRON-PRON-NC	84.2	55.8	67.1	88.3	68.7	77.3
ENT-PRON	78.8	68.7	73.4	80.4	69.8	74.7
MATCH	85.6	90.2	87.8	85.3	93.7	89.3
CONTAINS	72.4	80.9	76.4	77.4	78.9	78.1
OTHER	60.1	70.2	64.7	71.7	64.2	67.7

Table 4: Pairwise performance by category, on the test set of the English OntoNotes 5.0 dataset. LINGMESS surpasses the *s2e* model (Kirstain et al., 2021) for most categories by a substantial margin.

Inference At inference time, we compute the score of each mention based on the shared scorer as well as the per-type scorer matching the mention type. We then form the coreference chains by linking each mention q to its most likely antecedent c according to $F(c, q)$. We do not use higher-order inference as it has been shown to have a marginal impact (Xu and Choi, 2020).

4 Experiments

Our baseline is the *s2e* model trained on the English OntoNotes 5.0 dataset by its authors (Kirstain et al., 2021). We train LINGMESS also on OntoNotes, and evaluate both models on OntoNotes, WikiCoref, GAP, WinoGender, WinoBias and BUG. Implementation details are described in Appendix B.

Performance Table 2 presents the performance of LINGMESS on the test set of Ontonotes. LINGMESS achieves 81.4 F1 on Ontonotes, while

the *s2e* baseline achieves 80.3. Our performance gain is statistically significant according to a non-parametric permutation test ($p < 0.05$). Additionally, Table 3 shows that LINGMESS outperforms the *s2e* model on WikiCoref (+2.9) GAP (+1.3), WinoGender (+6.8), WinoBias (+0.8) and BUG (+2.4), indicating a better out-of-domain generalization.

Importance of per-category scoring. To assess that the improvement of LINGMESS is due to the decomposition into our set of categories and not to the added parameters, we do two experiments. First, we train a random baseline, which randomly assigns a category for each pair⁷ and obtain similar results as the baseline. Second, we train our model by optimizing only the overall loss L_{COREF} and not $L_{EXPERTS}$. This achieves lower results than the baseline, due to low mention recall.

In addition to the standard coreference evaluation, we report pairwise performance for each category. Given a mention-pair (c, q) , if $F(c, q)$ is higher than 0, we treat it as a positive prediction, otherwise negative. We then measure precision, recall and F1 based on gold clusters labels. Table 4 shows the pairwise performance of the *s2e* model and LINGMESS. LINGMESS outperforms *s2e* by a significant margin for all categories (e.g +7.4 F1 for PRON-PRON-C, +10.2 F1 for PRON-PRON-NC, etc.).⁸ The performance varies across the different categories, suggesting aspects of the coreference problem where future work can attempt to improve.

The importance of the shared scorer. To investigate the role of the shared scorer, we trained the LINGMESS model with only the per-type pairwise scorers, excluding the shared pairwise scorer $F_s(c, q)$ and its accompanying loss term $L_s(q)$.

⁷For each pair of mentions (c, q) , we take the modulo of the sum of the ASCII code of the last character of the last token of c and q .

⁸These gains in this pairwise metric are higher than the CoNLL metrics reported in Table 2, because the CoNLL metrics are based on the final clusters, after aggregation of individual pairwise decisions.

This resulted in a significant decrease in performance (-0.9), specifically in the recall of the mention detection component. However, adding the shared scorer was able to mitigate this degradation by balancing the different “experts” pairwise scorers.

5 Related Work

Many pre-neural works consider the various linguistic phenomena involved in coreference resolution as a different challenge. The early coreference system by Zhou and Su (2004) divided the antecedents candidates into distinct coreference categories (e.g., Name Alias, Apposition, Definite Noun, and a few more) and defined tailored rules for each category. Later, Lee et al. (2013) proposed the multi-sieves deterministic model, where each sieve adds coreference links between mention pairs from a specific linguistic category (e.g string match, compatible pronoun, etc.). Haghghi and Klein (2009) performed an error analysis of their coreference model according to different types of antecedent decisions, such as Proper Noun-Pronoun, Pronoun-Pronoun, etc. Based on this analysis, they focus on fixing the pronoun antecedent choices by adding syntactic features. More recently, Lu and Ng (2020) analyze empirically the performance of neural coreference resolvers on various fine-grained resolution categories of mentions (e.g gendered pronoun vs. 1st and 2nd pronoun). They find that while models perform well on name and nominal mention pairs with some shared content words, they still struggle with resolving pronouns, particularly relative pronouns.

Early supervised statistical models train a feature-based classifier that incorporates the type of antecedent decision (e.g. pronoun-entity, string match) as features at the mention-pair level (Soon et al., 2001; Bengtson and Roth, 2008; Clark and Manning, 2015, 2016). Subsequently, Denis and Baldridge (2008) demonstrate that training separate classifiers that specialize in particular types of mentions (e.g third person pronouns, speech pronouns, proper names, definite descriptions, and all other) provides significant performance improvements. Lassalle and Denis (2013) took that observation a step further and proposed a more advanced method for model specialization by learning to separate types of mention into optimal classes and their proper feature space.

In our work, we make progress in coreference

systems specialization direction, and show that the incorporation of linguistic information is helpful also in the context of end-to-end neural models.

6 Conclusion

We present LINGMESS, a coreference model that significantly improves accuracy by splitting the scoring function into different categories, and routing each scoring decision to its own category based on a deterministic, linguistically informed heuristic. This indicates that while end-to-end training is very effective, linguistic knowledge and symbolic computation can still be used to improve results.

Limitations

In this paper, we consider a set of 6 linguistic categories of mention pairs, as listed in Table 1. These categories might not be optimal for the task, while a different set of finer-grained categories may result to a higher performance gain. Another aspect that can be considered as a limitation is the computation of the categories for every possible pairs. Although the model considers only the top-scoring spans, this additional computation layer increases training and inference time over the baseline (see Appendix B.3 for the exact time). Our linguistic heuristics could be improved by, e.g., running a parser and considering head words. However, we chose not to do so in this work as this will further increase runtime.

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement No. 802774 (iEXTRACT). Arie Cattan is partially supported by the PBC fellowship for outstanding PhD candidates in data science.

References

- Amit Bagga and Breck Baldwin. 1998. [Entity-based cross-document coreferencing using the vector space model](#). In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 79–85, Montreal, Quebec, Canada. Association for Computational Linguistics.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *ArXiv*, abs/2004.05150.

- Eric Bengtson and Dan Roth. 2008. [Understanding the value of features for coreference resolution](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 294–303, Honolulu, Hawaii. Association for Computational Linguistics.
- Kevin Clark and Christopher D. Manning. 2015. [Entity-centric coreference resolution with model stacking](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1405–1415, Beijing, China. Association for Computational Linguistics.
- Kevin Clark and Christopher D. Manning. 2016. [Improving coreference resolution by learning entity-level distributed representations](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 643–653, Berlin, Germany. Association for Computational Linguistics.
- Pascal Denis and Jason Baldridge. 2008. [Specialized models and ranking for coreference resolution](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 660–669, Honolulu, Hawaii. Association for Computational Linguistics.
- Vladimir Dobrovolskii. 2021. [Word-level coreference resolution](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7670–7675, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Abbas Ghaddar and Phillippe Langlais. 2016. [WikiCoref: An English coreference-annotated corpus of Wikipedia articles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 136–142, Portorož, Slovenia. European Language Resources Association (ELRA).
- Aria Haghighi and Dan Klein. 2009. [Simple coreference resolution with rich syntactic and semantic features](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 1152–1161, Singapore. Association for Computational Linguistics.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Mandar Joshi, Omer Levy, Luke Zettlemoyer, and Daniel Weld. 2019. [BERT for coreference resolution: Baselines and analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5803–5808, Hong Kong, China. Association for Computational Linguistics.
- Ben Kantor and Amir Globerson. 2019. [Coreference resolution with entity equalization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 673–677, Florence, Italy. 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.
- Emmanuel Lassalle and Pascal Denis. 2013. [Improving pairwise coreference models through feature space hierarchy learning](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 497–506, Sofia, Bulgaria. Association for Computational Linguistics.
- Heeyoung Lee, Angel Chang, Yves Peirsman, Nathanael Chambers, Mihai Surdeanu, and Dan Jurafsky. 2013. [Deterministic coreference resolution based on entity-centric, precision-ranked rules](#). *Computational Linguistics*, 39(4):885–916.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. [End-to-end neural coreference resolution](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197, Copenhagen, Denmark. Association for Computational Linguistics.
- Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. [Higher-order coreference resolution with coarse-to-fine inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 687–692, New Orleans, Louisiana. Association for Computational Linguistics.
- Shahar Levy, Koren Lazar, and Gabriel Stanovsky. 2021. [Collecting a large-scale gender bias dataset for coreference resolution and machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2470–2480, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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.
- Xiaoqiang Luo. 2005. [On coreference resolution performance metrics](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages

- 25–32, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Nafise Sadat Moosavi and Michael Strube. 2016. [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)*, pages 632–642, Berlin, Germany. Association for Computational Linguistics.
- Shon Otmazgin, Arie Cattan, and Yoav Goldberg. 2022. [F-coref: Fast, accurate and easy to use coreference resolution.](#) In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 48–56, Taipei, Taiwan. Association for Computational Linguistics.
- 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.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. [Gender bias in coreference resolution.](#) 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 8–14, New Orleans, Louisiana. Association for Computational Linguistics.
- Wee Meng Soon, Hwee Tou Ng, and Daniel Chung Yong Lim. 2001. [A machine learning approach to coreference resolution of noun phrases.](#) *Computational Linguistics*, 27(4):521–544.
- Raghuveer Thirukovalluru, Nicholas Monath, Kumar Shridhar, Manzil Zaheer, Mrinmaya Sachan, and Andrew McCallum. 2021. [Scaling within document coreference to long texts.](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3921–3931, Online. Association for Computational Linguistics.
- Shubham Toshniwal, Sam Wiseman, Allyson Ettinger, Karen Livescu, and Kevin Gimpel. 2020. [Learning to Ignore: Long Document Coreference with Bounded Memory Neural Networks.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8519–8526, Online. Association for Computational Linguistics.
- Dennis Ulmer, Christian Hardmeier, and Jes Frellsen. 2022. [deep-significance-easy and meaningful statistical significance testing in the age of neural networks.](#) *arXiv preprint arXiv:2204.06815*.
- Marc Vilain, John Burger, John Aberdeen, Dennis Connolly, and Lynette Hirschman. 1995. [A model-theoretic coreference scoring scheme.](#) In *Sixth Message Understanding Conference (MUC-6): Proceedings of a Conference Held in Columbia, Maryland, November 6-8, 1995*.
- 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.
- 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.
- Patrick Xia, João Sedoc, and Benjamin Van Durme. 2020. [Incremental neural coreference resolution in constant memory.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8617–8624, Online. Association for Computational Linguistics.
- Liyang Xu and Jinho D. Choi. 2020. [Revealing the myth of higher-order inference in coreference resolution.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8527–8533, Online. Association for Computational Linguistics.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. [Gender bias in coreference resolution: Evaluation and debiasing methods.](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 15–20, New

Orleans, Louisiana. Association for Computational Linguistics.

GuoDong Zhou and Jian Su. 2004. [A high-performance coreference resolution system using a constraint-based multi-agent strategy](#). In *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics*, pages 522–528, Geneva, Switzerland. COLING.

A Model Architecture

Given a sequence of tokens x_1, \dots, x_n from an input document, a transformer-based (BERT-like) encoder first forms contextualized representation vectors, $\mathbf{x}_1, \dots, \mathbf{x}_n$ for each token in the sequence.

A.1 The *s2e* Model

Mention scorer Given a span $q = (\mathbf{x}_i, \mathbf{x}_j)$, represented by its start and end tokens, the score for q being a mention is defined as follow:

$$\mathbf{m}_s(\mathbf{x}) = \text{GeLU}(\mathbf{W}_{m_s}\mathbf{x}) \quad \mathbf{m}_e(\mathbf{x}) = \text{GeLU}(\mathbf{W}_{m_e}\mathbf{x})$$

$$\begin{aligned} f_m(q) &= \mathbf{m}_s(\mathbf{x}_i) \cdot \mathbf{v}_s + \mathbf{m}_s(\mathbf{x}_j) \cdot \mathbf{v}_e \\ &\quad + \mathbf{m}_s(\mathbf{x}_i) \cdot \mathbf{B}_m \cdot \mathbf{m}_s(\mathbf{x}_j) \end{aligned}$$

where $\mathbf{m}_s(\mathbf{x})$ and $\mathbf{m}_e(\mathbf{x})$ are two non-linear functions to obtain start and end representations for each token x , and $f_m(q)$ is a biaffine product over these representations.

Antecedent scorer Given two spans, $c = (\mathbf{x}_i, \mathbf{x}_j)$ and $q = (\mathbf{x}_k, \mathbf{x}_l)$, represented by their start and end tokens, the score for c being an antecedent of q is computed as follow:

$$\mathbf{a}_s(\mathbf{x}) = \text{GeLU}(\mathbf{W}_{a_s}\mathbf{x}) \quad \mathbf{a}_e(\mathbf{x}) = \text{GeLU}(\mathbf{W}_{a_e}\mathbf{x})$$

$$\begin{aligned} f_a(c, q) &= \mathbf{a}_s(\mathbf{x}_i) \cdot \mathbf{B}_{ss} \cdot \mathbf{a}_s(\mathbf{x}_k) \\ &\quad + \mathbf{a}_e(\mathbf{x}_j) \cdot \mathbf{B}_{es} \cdot \mathbf{a}_s(\mathbf{x}_k) \\ &\quad + \mathbf{a}_s(\mathbf{x}_i) \cdot \mathbf{B}_{se} \cdot \mathbf{a}_e(\mathbf{x}_l) \\ &\quad + \mathbf{a}_e(\mathbf{x}_j) \cdot \mathbf{B}_{ee} \cdot \mathbf{a}_e(\mathbf{x}_l) \end{aligned}$$

Similar to the mention scorer, $\mathbf{a}_s(\mathbf{x})$ and $\mathbf{a}_e(\mathbf{x})$ are two non-linear functions to obtain start/end representations for each token, and $f_a(c, q)$ is a sum of four bilinear functions over the start and end representations of c and q .

A.2 LINGMESS

Mention scorer Our mention scorer is the same as *s2e* mention scorer implementation.

Antecedent scorer As mentioned in the paper (§3), in addition to the shared antecedent scorer $f_a(c, q)$, LINGMESS includes a dedicated antecedent scorer $f_a^t(c, q)$ for each category $t \in \{\text{PRON-PRON-C, PRON-PRON-NC, ENT-PRON, MATCH, CONTAINS, OTHER}\}$. The overall score for $c = (\mathbf{x}_i, \mathbf{x}_j)$ being an antecedent of $q = (\mathbf{x}_k, \mathbf{x}_l)$ becomes the sum of the shared scorer and the relevant category “expert” scorer:

$$f(c, q) = f_a(c, q) + f_a^{T(c, q)}(c, q)$$

where $T(c, q)$ is a deterministic function to determine the category t of the pair (c, q) .

For each category t , we define two specific non-linear functions to obtain start and end representations ($\mathbf{a}_s^t(\mathbf{x})$ and $\mathbf{a}_e^t(\mathbf{x})$) as well as an “expert” antecedent scoring function $f_a^t(c, q)$:

$$\mathbf{a}_s^t(\mathbf{x}) = \text{GeLU}(\mathbf{W}_{a_s^t}\mathbf{x}) \quad \mathbf{a}_e^t(\mathbf{x}) = \text{GeLU}(\mathbf{W}_{a_e^t}\mathbf{x})$$

$$\begin{aligned} f_a^t(c, q) &= \mathbf{a}_s^t(\mathbf{x}_i) \cdot \mathbf{B}_{ss}^t \cdot \mathbf{a}_s^t(\mathbf{x}_k) \\ &\quad + \mathbf{a}_e^t(\mathbf{x}_j) \cdot \mathbf{B}_{es}^t \cdot \mathbf{a}_s^t(\mathbf{x}_k) \\ &\quad + \mathbf{a}_s^t(\mathbf{x}_i) \cdot \mathbf{B}_{se}^t \cdot \mathbf{a}_e^t(\mathbf{x}_l) \\ &\quad + \mathbf{a}_e^t(\mathbf{x}_j) \cdot \mathbf{B}_{ee}^t \cdot \mathbf{a}_e^t(\mathbf{x}_l) \end{aligned}$$

Overall, our model introduces 6 learnable matrices for each category ($\mathbf{W}_{a_s^t}^t, \mathbf{W}_{a_e^t}^t, \mathbf{B}_{ss}^t, \mathbf{B}_{es}^t, \mathbf{B}_{se}^t, \mathbf{B}_{ee}^t$). The transformer-based encoder and the mention scorer are shared between all the different pairwise scorers.

B Implementation Details

B.1 Hyperparameters

We extend the *s2e*’s implementation based on PyTorch (Paszke et al., 2019) and Transformers (Wolf et al., 2020). We used the same hyperparameters (e.g learning rate, warmup, etc.) as the *s2e* model except the hidden size of all matrices W and B . As our method introduces a dedicated antecedent scoring function f_a^t function for each category t , we reduce the size of these matrices from 3072 to 2048 to fit training into memory in our hardware. Similar to the baseline our head method is on top of Longformer-Large (Beltagy et al., 2020), resulting in a total of 590M learnable parameters (the *s2e* model contains 494M learnable parameters). We used dynamic batching both for training and inference, specifically 5K tokens in a single batch during training and 10K tokens at inference.

	Kirstain et al. (2021)			LINGMESS		
	P	R	F1	P	R	F1
PRON-PRON-C	91.7	77.5	84.0	91.7	90.2	91.0
PRON-PRON-NC	88.9	66.2	75.9	90.2	81.3	85.5
ENT-PRON	82.0	74.1	77.9	81.4	74.7	77.9
MATCH	88.3	87.5	87.9	88.4	92.0	90.2
CONTAINS	69.1	77.2	72.9	76.1	73.5	74.8
OTHER	56.8	67.5	61.7	70.8	64.4	67.5

Table 5: Pairwise performance by category, on the dev set of the English OntoNotes 5.0 dataset.

	Masc	Fem	Bias	Overall
Kirstain et al. (2021)	90.6	85.8	0.95	88.3
LINGMESS	91.3	87.8	0.96	89.6

Table 6: Performance on the test set of the GAP coreference dataset. The reported metrics are F1 scores.

B.2 Evaluation

As mentioned in the paper (§4), we conduct our experiments on the English portion of the OntoNotes corpus (Pradhan et al., 2012). This dataset contains 2802 documents for training, 343 for development, and 348 for test.

We evaluate our model according to the standard coreference metric: MUC (Vilain et al., 1995), B³ (Bagga and Baldwin, 1998), CEAF_{φ₄} (Luo, 2005), and LEA (Moosavi and Strube, 2016) using the official CoNLL coreference scorer.⁹ LINGMESS achieves 81.6 CoNLL F1 on the development set of Ontonotes. Also, Table 5 presents the pairwise performance on the development set for each category. We compute statistical significance with a non-parametric permutation test using Ulmer et al. (2022)’s implementation. Table 6 shows that LINGMESS consistently outperforms the *s2e* model on GAP.

B.3 Runtime and Memory

Our model was trained for 129 epochs on a single 32GB NVIDIA Tesla V100 SXM2 GPU. The training took 23 hours. As shown in Table 7, the run-time at inference time in LINGMESS is longer than in the *s2e* model because of the category selection for every possible pair of mentions. The memory consumption remains quite similar to the baseline.

⁹<https://github.com/conll/reference-coreference-scorers>.

	Runtime	Memory
Kirstain et al. (2021)	28	5.4
LINGMESS	43	5.9

Table 7: Inference time(Seconds) and memory(GiB) on 343 docs of OntoNotes development set. Using Dynamic batching, 10K tokens in a single batch. Hardware, NVIDIA Tesla V100 SXM2

ID	Pronouns
1	<i>I, me, my, mine, myself</i>
2	<i>you, your, yours, yourself, yourselves</i>
3	<i>he, him, his, himself</i>
4	<i>she, her, hers, herself</i>
5	<i>it, its, itself</i>
6	<i>we, us, our, ours, ourselves</i>
7	<i>they, them, their, themselves</i>
8	<i>that, this</i>

Table 8: List of groups of compatible pronouns, pronouns with the same ID are considered as compatible.

C Determining pair types

Our method routes each pair of spans to their corresponding category scorer. This decision is based on the linguistic properties of the spans. Given a mention-pair (c, q) , we defined a rule based function $T(c, q)$ that determines the category of that pair. If c and q are both pronouns, if they are compatible according to gender, number and animacy (see Table 8 for the full list), the mention pair will be routed to PRON-PRON-C, otherwise PRON-PRON-NC. If c is pronoun and q is a non-pronoun span (or vice-versa) we route the mention-pair to PRON-ENT. We route the remaining pairs to their corresponding categories (MATCH, CONTAINS or OTHER) by considering only content words, excluding the following stop words: $\{s, a, all, an, and, at, for, from, in, into, more, of, on, or, some, the, these, those\}$. Accordingly, the mentions “*the U.S. and Japan*” and “*Japan and the U.S.*” are considered MATCH, “*This lake of fire*” and “*the lake of fire*” are considered CONTAINS and “*Bill Clinton*” and “*The President*” are considered OTHER.