

BERT for Coreference Resolution: Baselines and Analysis

Mandar Joshi[†] Omer Levy[§] Daniel S. Weld^{†ε} Luke Zettlemoyer^{†§}

[†] Allen School of Computer Science & Engineering, University of Washington, Seattle, WA
{mandar90, weld, lsz}@cs.washington.edu

^εAllen Institute of Artificial Intelligence, Seattle
{danw}@allenai.org

[§]Facebook AI Research, Seattle
{omerlevy, lsz}@fb.com

Abstract

We apply BERT to coreference resolution, achieving strong improvements on the OntoNotes (+3.9 F1) and GAP (+11.5 F1) benchmarks. A qualitative analysis of model predictions indicates that, compared to ELMo and BERT-base, BERT-large is particularly better at distinguishing between related but distinct entities (e.g., President and CEO). However, there is still room for improvement in modeling document-level context, conversations, and mention paraphrasing. Our code and models are publicly available¹.

1 Introduction

Recent BERT-based models have reported dramatic gains on multiple semantic benchmarks including question-answering, natural language inference, and named entity recognition (Devlin et al., 2019). Apart from better bidirectional reasoning, one of BERT’s major improvements over previous methods (Peters et al., 2018; McCann et al., 2017) is passage-level training,² which allows it to better model longer sequences.

We fine-tune BERT to coreference resolution, achieving strong improvements on the GAP (Webster et al., 2018) and OntoNotes (Pradhan et al., 2012) benchmarks. We present two ways of extending the c2f-coref model in Lee et al. (2018). The *independent* variant uses non-overlapping segments each of which acts as an independent instance for BERT. The *overlap* variant splits the document into overlapping segments so as to provide the model with context beyond 512 tokens. BERT-large improves over ELMo-based c2f-coref 3.9% on OntoNotes and 11.5% on GAP (both absolute).

¹<https://github.com/mandarjosshi90/coref>

²Each BERT training example consists of around 512 word pieces, while ELMo is trained on single sentences.

A qualitative analysis of BERT and ELMo-based models (Table 3) suggests that BERT-large (unlike BERT-base) is remarkably better at distinguishing between related yet distinct entities or concepts (e.g., Repulse Bay and Victoria Harbor). However, both models often struggle to resolve coreferences for cases that require world knowledge (e.g., *the developing story* and *the scandal*). Likewise, modeling pronouns remains difficult, especially in conversations.

We also find that BERT-large benefits from using longer context windows (384 word pieces) while BERT-base performs better with shorter contexts (128 word pieces). Yet, both variants perform much worse with longer context windows (512 tokens) in spite of being trained on 512-size contexts. Moreover, the *overlap* variant, which artificially extends the context window beyond 512 tokens provides no improvement. This indicates that using larger context windows for pretraining might not translate into effective long-range features for a downstream task. Larger models also exacerbate the memory-intensive nature of span representations (Lee et al., 2017), which have driven recent improvements in coreference resolution. Together, these observations suggest that there is still room for improvement in modeling document-level context, conversations, and mention paraphrasing.

2 Method

For our experiments, we use the higher-order coreference model in Lee et al. (2018) which is the current state of the art for the English OntoNotes dataset (Pradhan et al., 2012). We refer to this as *c2f-coref* in the paper.

2.1 Overview of c2f-coref

For each mention span x , the model learns a distribution $P(\cdot)$ over possible antecedent spans Y :

$$P(y) = \frac{e^{s(x,y)}}{\sum_{y' \in Y} e^{s(x,y')}} \quad (1)$$

The scoring function $s(x, y)$ between spans x and y uses fixed-length span representations, \mathbf{g}_x and \mathbf{g}_y to represent its inputs. These consist of a concatenation of three vectors: the two LSTM states of the span endpoints and an attention vector computed over the span tokens. It computes the score $s(x, y)$ by the mention score of x (i.e. how likely is the span x to be a mention), the mention score of y , and the joint compatibility score of x and y (i.e. assuming they are both mentions, how likely are x and y to refer to the same entity). The components are computed as follows:

$$s(x, y) = s_m(x) + s_m(y) + s_c(x, y) \quad (2)$$

$$s_m(x) = \text{FFNN}_m(\mathbf{g}_x) \quad (3)$$

$$s_c(x, y) = \text{FFNN}_c(\mathbf{g}_x, \mathbf{g}_y, \phi(x, y)) \quad (4)$$

where $\text{FFNN}(\cdot)$ represents a feedforward neural network and $\phi(x, y)$ represents speaker and meta-data features. These span representations are later refined using antecedent distribution from a span-ranking architecture as an attention mechanism.

2.2 Applying BERT

We replace the entire LSTM-based encoder (with ELMo and GloVe embeddings as input) in c2f-coref with the BERT transformer. We treat the first and last word-pieces (concatenated with the attended version of all word pieces in the span) as span representations. Documents are split into segments of `max_segment_len`, which we treat as a hyperparameter. We experiment with two variants of splitting:

Independent The *independent* variant uses non-overlapping segments each of which acts as an independent instance for BERT. The representation for each token is limited to the set of words that lie in its segment. As BERT is trained on sequences of at most 512 word pieces, this variant has limited encoding capacity especially for tokens that lie at the start or end of their segments.

Overlap The *overlap* variant splits the document into overlapping segments by creating a T -sized segment after every $T/2$ tokens. These segments are then passed on to the BERT encoder independently, and the final token representation is derived by element-wise interpolation of representations from both overlapping segments.

Let $\mathbf{r}_1 \in \mathbb{R}^d$ and $\mathbf{r}_2 \in \mathbb{R}^d$ be the token representations from the overlapping BERT segments. The final representation $\mathbf{r} \in \mathbb{R}^d$ is given by:

$$\mathbf{f} = \sigma(\mathbf{w}^T [\mathbf{r}_1; \mathbf{r}_2]) \quad (5)$$

$$\mathbf{r} = \mathbf{f} \cdot \mathbf{r}_1 + (1 - \mathbf{f}) \cdot \mathbf{r}_2 \quad (6)$$

where $\mathbf{w} \in \mathbb{R}^{2d \times d}$ is a trained parameter and $[\cdot; \cdot]$ represents concatenation. This variant allows the model to artificially increase the context window beyond the `max_segment_len` hyperparameter.

All layers in both model variants are then fine-tuned following Devlin et al. (2019).

3 Experiments

We evaluate our BERT-based models on two benchmarks: the paragraph-level GAP dataset (Webster et al., 2018), and the document-level English OntoNotes 5.0 dataset (Pradhan et al., 2012). OntoNotes examples are considerably longer and typically require multiple segments to read the entire document.

Implementation and Hyperparameters We extend the original Tensorflow implementations of c2f-coref³ and BERT.⁴ We fine tune all models on the OntoNotes English data for 20 epochs using a dropout of 0.3, and learning rates of 1×10^{-5} and 2×10^{-4} with linear decay for the BERT parameters and the task parameters respectively. We found that this made a sizable impact of 2-3% over using the same learning rate for all parameters.

We trained separate models with `max_segment_len` of 128, 256, 384, and 512; the models trained on 128 and 384 word pieces performed the best for BERT-base and BERT-large respectively. As span representations are memory intensive, we truncate documents randomly to eleven segments for BERT-base and 3 for BERT-large during training. Likewise, we use a batch size of 1 document following (Lee et al., 2018). While training the large model requires 32GB GPUs, all models can be tested on 16GB GPUs. We use the cased English variants in all our experiments.

Baselines We compare the c2f-coref + BERT system with two main baselines: (1) the original ELMo-based c2f-coref system (Lee et al., 2018), and (2) its predecessor, e2e-coref (Lee et al.,

³<http://github.com/kenton1/e2e-coref/>

⁴<https://github.com/google-research/bert>

	MUC			B ³			CEAF _{φ₄}			Avg. F1
	P	R	F1	P	R	F1	P	R	F1	
Martschat and Strube (2015)	76.7	68.1	72.2	66.1	54.2	59.6	59.5	52.3	55.7	62.5
(Clark and Manning, 2015)	76.1	69.4	72.6	65.6	56.0	60.4	59.4	53.0	56.0	63.0
(Wiseman et al., 2015)	76.2	69.3	72.6	66.2	55.8	60.5	59.4	54.9	57.1	63.4
Wiseman et al. (2016)	77.5	69.8	73.4	66.8	57.0	61.5	62.1	53.9	57.7	64.2
Clark and Manning (2016)	79.2	70.4	74.6	69.9	58.0	63.4	63.5	55.5	59.2	65.7
e2e-coref (Lee et al., 2017)	78.4	73.4	75.8	68.6	61.8	65.0	62.7	59.0	60.8	67.2
c2f-coref (Lee et al., 2018)	81.4	79.5	80.4	72.2	69.5	70.8	68.2	67.1	67.6	73.0
Fei et al. (2019)	85.4	77.9	81.4	77.9	66.4	71.7	70.6	66.3	68.4	73.8
EE (Kantor and Globerson, 2019)	82.6	84.1	83.4	73.3	76.2	74.7	72.4	71.1	71.8	76.6
BERT-base + c2f-coref (independent)	80.2	82.4	81.3	69.6	73.8	71.6	69.0	68.6	68.8	73.9
BERT-base + c2f-coref (overlap)	80.4	82.3	81.4	69.6	73.8	71.7	69.0	68.5	68.8	73.9
BERT-large + c2f-coref (independent)	84.7	82.4	83.5	76.5	74.0	75.3	74.1	69.8	71.9	76.9
BERT-large + c2f-coref (overlap)	85.1	80.5	82.8	77.5	70.9	74.1	73.8	69.3	71.5	76.1

Table 1: OntoNotes: BERT improves the c2f-coref model on English by 0.9% and 3.9% respectively for base and large variants. The main evaluation is the average F1 of three metrics – MUC, B³, and CEAF_{φ₄} on the test set.

Model	M	F	B	O
e2e-coref	67.2	62.2	0.92	64.7
c2f-coref	75.8	71.1	0.94	73.5
BERT + RR Liu et al. (2019)	80.3	77.4	0.96	78.8
BERT-base + c2f-coref	84.4	81.2	0.96	82.8
BERT-large + c2f-coref	86.9	83.0	0.95	85.0

Table 2: GAP: BERT improves the c2f-coref model by 11.5%. The metrics are F1 score on Masculine and Feminine examples, Overall, and a Bias factor (F / M).

2017), which does not use contextualized representations. In addition to being more computationally efficient than e2e-coref, c2f-coref iteratively refines span representations using attention for higher-order reasoning.

3.1 Paragraph Level: GAP

GAP (Webster et al., 2018) is a human-labeled corpus of ambiguous pronoun-name pairs derived from Wikipedia snippets. Examples in the GAP dataset fit within a single BERT segment, thus eliminating the need for cross-segment inference. Following Webster et al. (2018), we trained our BERT-based c2f-coref model on OntoNotes.⁵ The predicted clusters were scored against GAP examples according to the official evaluation script. Table 2 shows that BERT improves c2f-coref by 9% and 11.5% for the base and large models respectively. These results are in line with large gains reported for a variety of semantic tasks by BERT-based models (Devlin et al., 2019).

⁵This is motivated by the fact that GAP, with only 4,000 name-pronoun pairs in its dev set, is not intended for full-scale training.

3.2 Document Level: OntoNotes

OntoNotes (English) is a document-level dataset from the CoNLL-2012 shared task on coreference resolution. It consists of about one million words of newswire, magazine articles, broadcast news, broadcast conversations, web data and conversational speech data, and the New Testament. The main evaluation is the average F1 of three metrics – MUC, B³, and CEAF_{φ₄} on the test set according to the official CoNLL-2012 evaluation scripts.

Table 1 shows that BERT-base offers an improvement of 0.9% over the ELMo-based c2f-coref model. Given how gains on coreference resolution have been hard to come by as evidenced by the table, this is still a considerable improvement. However, the magnitude of gains is relatively modest considering BERT’s arguably better architecture and many more trainable parameters. This is in sharp contrast to how even the base variant of BERT has very substantially improved the state of the art in other tasks. BERT-large, however, improves c2f-coref by the much larger margin of 3.9%. We also observe that the *overlap* variant offers no improvement over *independent*.

Concurrent with our work, Kantor and Globerson (2019), who use higher-order entity-level representations over “frozen” BERT features, also report large gains over c2f-coref. While their feature-based approach is more memory efficient, the fine-tuned model seems to yield better results. Also concurrent, SpanBERT (Joshi et al., 2019), another self-supervised method, pretrains span representations achieving state of the art results (Avg. F1 79.6) with the *independent* variant.

Category	Snippet	#base	#large
Related Entities	Watch spectacular performances by dolphins and sea lions at the <i>Ocean Theater</i> ... It seems the North Pole and the <u>Marine Life Center</u> will also be renovated.	12	7
Lexical	Over the past 28 years , <i>the Ocean Park</i> has basically.. The entire park has been ...	15	9
Pronouns	In the meantime , our children need <i>an education</i> . That's all we're asking.	17	13
Mention Paraphrasing	And in case you missed it <i>the Royals</i> are here. Today Britain's Prince Charles and his wife Camilla ...	14	12
Conversation	(Priscilla:) My mother was Thelma Wahl . She was ninety years old ... (Keith:) <i>Priscilla Scott</i> is mourning . <i>Her</i> mother Thelma Wahl was a resident ..	18	16
Misc.	He is my , She is my Goddess , ah	17	17
<i>Total</i>		93	74

Table 3: Qualitative Analysis: #base and #large refers to the number of cluster-level errors on a subset of the OntoNotes English development set. Underlined and **bold-faced** mentions respectively indicate incorrect and missing assignments to *italicized* mentions/clusters. The miscellaneous category refers to other errors including (reasonable) predictions that are either missing from the gold data or violate annotation guidelines.

Doc length	#Docs	Spread	F1 (base)	F1 (large)
0 - 128	48	37.3	80.6	84.5
128 - 256	54	71.7	80.0	83.0
256 - 512	74	109.9	78.2	80.0
512 - 768	64	155.3	76.8	80.2
768 - 1152	61	197.6	71.1	76.2
1152+	42	255.9	69.9	72.8
All	343	179.1	74.3	77.3

Table 4: Performance on the English OntoNotes dev set generally drops as the document length increases. Spread is measured as the average number of tokens between the first and last mentions in a cluster.

Segment Length	F1 (BERT-base)	F1 (BERT-large)
128	74.4	76.6
256	73.9	76.9
384	73.4	77.3
450	72.2	75.3
512	70.7	73.6

Table 5: Performance on the English OntoNotes dev set with varying values for `max_segment_len`. Neither model is able to effectively exploit larger segments; they perform especially badly when `maximum_segment_len` of 512 is used.

4 Analysis

We performed a qualitative comparison of ELMo and BERT models (Table 3) on the OntoNotes English development set by manually assigning error categories (e.g., pronouns, mention paraphrasing) to incorrect predicted clusters.⁶ Overall, we found 93 errors for BERT-base and 74 for BERT-large from the same 15 documents.

⁶Each incorrect cluster can belong to multiple categories.

Strengths We did not find salient qualitative differences between ELMo and BERT-base models, which is consistent with the quantitative results (Table 1). BERT-large improves over BERT-base in a variety of ways including pronoun resolution and lexical matching (e.g., *race track* and *track*). In particular, the BERT-large variant is better at distinguishing related, but distinct, entities. Table 3 shows several examples where the BERT-base variant merges distinct entities (like *Ocean Theater* and *Marine Life Center*) into a single cluster. BERT-large seems to be able to avoid such merging on a more regular basis.

Weaknesses An analysis of errors on the OntoNotes English development set suggests that better modeling of document-level context, conversations, and entity paraphrasing might further improve the state of the art.

Longer documents in OntoNotes generally contain larger and more spread-out clusters. We focus on three observations – (a) Table 4 shows how models perform distinctly worse on longer documents, (b) both models are unable to use larger segments more effectively (Table 5) and perform worse when the `max_segment_len` of 450 and 512 are used, and, (c) using overlapping segments to provide additional context does not improve results (Table 1). Recent work (Joshi et al., 2019) suggests that BERT’s inability to use longer sequences effectively is likely a by-product pretraining on short sequences for a vast majority of updates.

Comparing preferred segment lengths for base and large variants of BERT indicates that larger

models might better encode longer contexts. However, larger models also exacerbate the memory-intensive nature of span representations,⁷ which have driven recent improvements in coreference resolution. These observations suggest that future research in pretraining methods should look at more *effectively* encoding document-level context using sparse representations (Child et al., 2019).

Modeling *pronouns* especially in the context of conversations (Table 3), continues to be difficult for all models, perhaps partly because c2f-coref does very little to model dialog structure of the document. Lastly, a considerable number of errors suggest that models are still unable to resolve cases requiring *mention paraphrasing*. For example, bridging *the Royals* with *Prince Charles and his wife Camilla* likely requires pretraining models to encode relations between entities, especially considering that such learning signal is rather sparse in the training set.

5 Related Work

Scoring span or mention pairs has perhaps been one of the most dominant paradigms in coreference resolution. The base coreference model used in this paper from Lee et al. (2018) belongs to this family of models (Ng and Cardie, 2002; Bengtson and Roth, 2008; Denis and Baldridge, 2008; Fernandes et al., 2012; Durrett and Klein, 2013; Wiseman et al., 2015; Clark and Manning, 2016; Lee et al., 2017).

More recently, advances in coreference resolution and other NLP tasks have been driven by unsupervised contextualized representations (Peters et al., 2018; Devlin et al., 2019; McCann et al., 2017; Joshi et al., 2019). Of these, BERT (Devlin et al., 2019) notably uses pretraining on passage-level sequences (in conjunction with a bidirectional masked language modeling objective) to more effectively model long-range dependencies. SpanBERT (Joshi et al., 2019) focuses on pretraining span representations achieving current state of the art results on OntoNotes with the *independent* variant.

References

Eric Bengtson and Dan Roth. 2008. [Understanding the value of features for coreference resolution](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP '08*,

pages 294–303, Stroudsburg, PA, USA. Association for Computational Linguistics.

Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. [Generating long sequences with sparse transformers](#). *arXiv preprint arXiv:1904.10509*.

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. Association for Computational Linguistics.

Kevin Clark and Christopher D. Manning. 2016. [Deep reinforcement learning for mention-ranking coreference models](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2256–2262. 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. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Greg Durrett and Dan Klein. 2013. [Easy victories and uphill battles in coreference resolution](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1971–1982. Association for Computational Linguistics.

Hongliang Fei, Xu Li, Dingcheng Li, and Ping Li. 2019. [End-to-end deep reinforcement learning based coreference resolution](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 660–665, Florence, Italy. Association for Computational Linguistics.

Eraldo Fernandes, Cícero dos Santos, and Ruy Milidiú. 2012. [Latent structure perceptron with feature induction for unrestricted coreference resolution](#). In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 41–48. Association for Computational Linguistics.

Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2019. [SpanBERT: Improving pre-training by representing and predicting spans](#). *arXiv preprint arXiv:1907.10529*.

⁷We required a 32GB GPU to finetune BERT-large.

- 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.
- 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.
- Fei Liu, Luke Zettlemoyer, and Jacob Eisenstein. 2019. [The referential reader: A recurrent entity network for anaphora resolution](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5918–5925, Florence, Italy.
- Sebastian Martschat and Michael Strube. 2015. [Latent structures for coreference resolution](#). *Transactions of the Association for Computational Linguistics*, 3:405–418.
- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. 2017. [Learned in translation: Contextualized word vectors](#). In *Advances in Neural Information Processing Systems*, pages 6297–6308.
- Vincent Ng and Claire Cardie. 2002. [Identifying anaphoric and non-anaphoric noun phrases to improve coreference resolution](#). In *Proceedings of the 19th International Conference on Computational Linguistics - Volume 1, COLING '02*, pages 1–7, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. [Conll-2012 shared task: Modeling multilingual unrestricted coreference in ontonotes](#). In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40. Association for Computational Linguistics.
- 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.
- Sam Wiseman, Alexander M. Rush, Stuart Shieber, and Jason Weston. 2015. [Learning anaphoricity and antecedent ranking features for coreference resolution](#). 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 1416–1426. Association for Computational Linguistics.
- Sam Wiseman, Alexander M. Rush, and Stuart M. Shieber. 2016. [Learning global features for coreference resolution](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 994–1004. Association for Computational Linguistics.