

Anatomy of OntoGUM—Adapting GUM to the OntoNotes Scheme to Evaluate Robustness of SOTA Coreference Algorithms

Yilun Zhu¹, Sameer Pradhan^{2,3}, and Amir Zeldes¹

¹Department of Linguistics, Georgetown University
²Linguistic Data Consortium, University of Pennsylvania
³cemantix.org

yz565@georgetown.edu, pradhan@cemantix.org, Amir.Zeldes@georgetown.edu

Abstract

SOTA coreference resolution produces increasingly impressive scores on the OntoNotes benchmark. However lack of comparable data following the same scheme for more genres makes it difficult to evaluate generalizability to open domain data. Zhu et al. (2021) introduced the creation of the OntoGUM corpus for evaluating generalizability of the latest neural LM-based end-to-end systems. This paper covers details of the mapping process which is a set of deterministic rules applied to the rich syntactic and discourse annotations manually annotated in the GUM corpus. Out-of-domain evaluation across 12 genres shows nearly 15-20% degradation for both deterministic and deep learning systems, indicating a lack of generalizability or covert overfitting in existing coreference resolution models.

1 Introduction

Coreference resolution is the task of grouping referring expressions that point to the same entity, such as noun phrases and the pronouns that refer to them. The task entails detecting correct mention or ‘markable’ boundaries and creating a link with previous mentions, or antecedents. A coreference chain is a series of decisions which groups the markables into clusters. As a key component in Natural Language Understanding (NLU), the task can benefit a series of downstream applications such as Entity Linking, Dialogue Systems, Machine Translation, Summarization, and more (Poesio et al., 2016).

In recent years, deep learning models have achieved high scores in coreference resolution. The end-to-end approach (Lee et al., 2017, 2018) jointly scoring mention detection and coreference resolution currently not only beats earlier rule-based and statistical methods but also outperforms other deep learning approaches (Wiseman et al., 2016; Clark and Manning, 2016a,b). Additionally, language models trained on billions of words significantly

improve performance by providing rich word and context-level information for classifiers (Lee et al., 2018; Joshi et al., 2019a,b).

However, scores on the identity coreference layer of the benchmark OntoNotes dataset (Pradhan et al., 2013) do not reflect the generalizability of these systems. Moosavi and Strube (2017) pointed out that lexicalized coreference resolution models, including neural models using word embeddings, face a covert overfitting problem because of a large overlap between the vocabulary of coreferring mentions in the OntoNotes training and evaluation sets. This suggests that higher scores on OntoNotes-test may not indicate a better solution to the coreference resolution task.

To investigate the generalization problem of neural models, several projects have tested other datasets consistent with the OntoNotes scheme. Moosavi and Strube (2018) conducted out-of-domain evaluation on WikiCoref (Ghaddar and Langlais, 2016), a small dataset employing the same coreference definitions. Results showed that neural models (with fixed embeddings) do not achieve comparable performance (16.8% degradation in score) to scores on OntoNotes. More recently, the e2e model using BERT (Joshi et al., 2019b) showed gains on the GAP corpus (Webster et al., 2018) using contextualized embeddings; however GAP only contains name-pronoun coreference, a very specific subset of coreference, and is limited in domain to the same single source – Wikipedia.

Though previous work has already identified the overfitting problem, it has three main shortcomings. First, the scale of out-of-domain evaluation has been small and homogeneous: WikiCoref only contains 30 documents with ~60K tokens, much smaller than the OntoNotes test set, and the single genre Wiki domain in both WikiCoref and GAP is arguably not very far from some OntoNotes materials. Second, pretrained LMs, e.g. BERT (Devlin

Genre	Documents	Tokens	Mentions	Proper	Pron.	Other	Clusters
<i>academic</i>	16	15,112	1,232	283	262	687	421
<i>bio</i>	20	17,963	2,312	934	796	582	487
<i>conversation</i>	5	5,701	1,027	40	728	259	176
<i>fiction</i>	18	16,312	2,740	259	1,700	781	469
<i>interview</i>	19	18,060	2,622	501	1,223	898	608
<i>news</i>	21	14,094	1,803	796	340	667	477
<i>reddit</i>	18	16,286	2,297	117	1,336	844	578
<i>speech</i>	5	4,834	601	171	245	185	134
<i>textbook</i>	5	5,379	466	108	165	193	133
<i>vlog</i>	5	5,189	882	22	600	260	149
<i>voyage</i>	17	14,967	1,339	564	300	475	348
<i>whow</i>	19	16,927	2,057	53	1,001	1,003	491
Total	168	150,824	19,378	3,848	8,696	68,34	4,471

Table 1: Genre-breakdown Statistics of OntoGUM.

et al., 2019), popularized after the WikiCoref paper, can learn better representations of markables and surrounding sentences. Aside from GAP, which targets a highly specific subtask, no study has investigated whether contextualized embeddings encounter the same overfitting problem identified by Moosavi and Strube. Third, previous work may underestimate the performance degradation on WikiCoref in particular due to bias: In Moosavi and Strube (2018), embeddings were also trained on Wikipedia themselves, potentially making it easier for the model to learn coreference relations in Wikipedia text, despite limitations in other genres.

In this paper, we explore the generalizability of existing coreference models on a new benchmark dataset, which we make freely available. Compared with work using WikiCoref and GAP, our contributions can be summarized as follows:

- We propose OntoGUM, the largest open, gold standard dataset consistent with OntoNotes, with 168 documents (\sim 150K tokens, 19,378 mentions, 4,471 coref chains) in 12 genres,¹ including conversational genres, which complement OntoNotes for training and evaluation.
- We show that the SOTA neural model with contextualized embeddings encounters nearly 15% performance degradation on OntoGUM, showing that the overfitting problem is not overcome by contextualized language models.
- We give a genre-by-genre analysis for two popular systems, revealing relative strengths and weaknesses of current approaches and

¹**Written:** News/Fiction/Bio/Academic/Forum/Travel/How-to/Textbook; **Spoken:** Interview/Political/Vlog/Conversation.

the range of easier/more difficult targets for coreference resolution.

2 Related Work

OntoNotes and similar corpora OntoNotes is a human-annotated corpus with documents annotated with multiple layers of linguistic information including syntax, propositions, named entities, word senses, and within-document coreference (Weischedel et al., 2011; Pradhan et al., 2013). It covers three languages—English, Chinese and Arabic. The English subcorpus has 3,493 documents and \sim 1.6 million words annotated for coreference. WikiCoref, which is annotated for anaphoric relations, has 30 documents from English Wikipedia (Ghaddar and Langlais, 2016), containing 7,955 mentions in 1,785 chains, following OntoNotes guidelines.

GUM The Georgetown University Multilayer (GUM) corpus (Zeldes, 2017) is an open-source corpus of richly annotated texts from 12 genres, including 168 documents and over 150K tokens. Though it originally contains more coreference phenomena than OntoNotes using more exhaustive guidelines, it also contains rich syntactic, semantic and discourse annotations which allow us to create the OntoGUM dataset described below. The syntactic annotations, which consist of manually annotated Universal Dependencies trees (de Marneffe et al., 2021), are particularly useful in analyzing and converting the GUM corpus. We also note that due to its smaller size (currently about 10% the size of the OntoNotes coreference dataset), it is not possible to train SOTA neural approaches directly on this dataset while maintaining strong performance.

Other corpora As mentioned above, GAP is a gender-balanced labeled corpus of ambiguous

pronoun-name pairs, used for out-of-domain evaluation but limited in coreferent types and genre. Several other comprehensive coreference datasets exist as well, such as ARRAU (Poesio et al., 2018) and PreCo (Chen et al., 2018), but these corpora cannot be used for out-of-domain evaluation because they do not follow the OntoNotes scheme. Their conversion has not been attempted to date.

Coreference resolution systems Prior to the introduction of deep learning systems, the coreference task was approached using deterministic linguistic rules (Lee et al., 2013; Recasens et al., 2013) and statistical approaches (Durrett and Klein, 2013, 2014). More recently, three neural models achieved SOTA performance on this task: 1) ranking the candidate mention pairs (Wiseman et al., 2015; Clark and Manning, 2016a), 2) modeling global features of entity clusters (Clark and Manning, 2015, 2016b; Wiseman et al., 2016), and 3) end-to-end (e2e) approaches with joint loss for mention detection and coreferent pair scoring (Lee et al., 2017, 2018; Fei et al., 2019). The e2e method has become the dominant one, gaining the best scores on OntoNotes. To investigate differences between deterministic and deep learning models on unseen data, we evaluate the two approaches on OntoGUM.

3 Dataset Conversion

GUM’s annotation scheme subsumes all markables and coreference chains annotated in OntoNotes, meaning we do not need human annotation to recognize additional mentions in the conversion process, though mention boundaries differ subtly (e.g. for appositions and verbal mentions). Since GUM has gold syntax trees, we were able to process the entire conversion automatically. Additionally, most coreference evaluations use gold speaker information in OntoNotes, which is available in GUM (for *fiction*, *reddit* and spoken data) and could be assembled automatically as well.

The conversion is divided into two parts: removing coreference relations not included in the OntoNotes scheme, and removing or adjusting markables. For coreference relation deletion, we cut chains by removing expletive cataphora, and identifying the definiteness of nominal markables, since indefinites cannot be anaphors in OntoNotes - this was done based on the lemma of determiners attached to the head token of the span (or lack of a determiner) and the POS tag of the head. In addition to modifying existing mention clusters, we

Order of conversion steps
> Remove bridging & cataphora
> Contract verbal spans
> Merge appositions
> Remove NN compounding
> Remove copula
> Remove nested entities
> Adjust chains by definiteness
> Remove singletons

Table 2: The order of the steps in the dataset conversion from GUM to OntoNotes scheme.

also remove particular coreference relations and mention spans, such as Noun-Noun compounding (only included in OntoNotes for proper-name modifiers), bridging anaphora, copula predicates, nested entities (‘i-within-i’= single mentions containing coreferring pronouns), and singletons, i.e. mentions that are not referred to again (all not included in OntoNotes). We note that singletons are removed as the final step, in order to catch singletons generated during the conversion process. We also contract verbal markable spans to their head verb, and merge appositive constructions, which are explicitly marked in GUM, into single mentions, following the OntoNotes guidelines. The order of conversion steps are shown in Table 2.

3.1 Coreference relations

Cataphora Cataphora encompasses pronominal elements, including demonstratives, e.g. *those*, which precede an occurrence of a non-pronominal element that occurs within the same utterance and resolves their discourse referent. GUM specifies cataphora in the coreference annotation while OntoNotes only annotates pronominal markables and discards the resolved non-pronominal elements.

Unlike other relations, cataphora ‘points forward’, i.e. it is resolved by finding a subsequent lexical phrase corresponding to an earlier underspecified pronoun. As in (1), the pronoun *it* (bolded in the box) is resolved by the subsequent markable within the same utterance.

- (1) *Before:* **it**’s good to be able to do well at the World Cup , to be placed , but **it** also means that you get a really good opportunity to know where you ’re at in that two year gap between the Paralympics .
- (2) *After:* **it**’s good to be able to do well

at the World Cup , to be placed , but it also means that you get a really good opportunity to know where you 're at in that two year gap between the Paralympics .

Because GUM annotates coreference types for each relation, We create a heuristic algorithm to process cataphora. As in (2), the expletive pronoun (dashed box) is removed from the cluster it originally belongs to, leaving it as a singleton. Other coreference relations remain the same in that cluster.

Definiteness A definite marker indicates the referent is identifiable in a given context while indefinite markers often function to introduce new entities not mentioned before. Following these tendencies, OntoNotes does not allow coreference relations between an indefinite nominal and any kind of antecedents. GUM, however, only considers whether the markables refer to the same entity or not. For example, the indefinite noun phrase *a farm several miles outside of town* in (3) can occur in the middle of the coreferring chain.

(3) *Before:* Rachel Rook took Carroll home to meet her parents two months after she first slept with him ... her parents lived on a farm several miles outside of town ... He knew that they never left the farm ...

(4) *After:* Rachel Rook took Carroll home to meet her parents two months after she first slept with him ... her parents lived on a farm several miles outside of town ... He knew that they never left the farm ...

We use the POS tags, dependency labels, and lemmas to distinguish definite and indefinite markables. A definite nominal markable satisfies one of the following cases: (1) it is a pronoun; (2) the head noun is possessed; (3) it is a proper noun; and (4) anything which has the dependency relation DET with a definite determiner lemma such as *the*, *that*, etc. The converted document looks like (4), where the indefinite nominal phrase and its original antecedent (dashed box) fall in different clusters and the indefinite markable is the first mention in the new cluster².

²If the markable *home* also has an antecedent, the coreference relation is not affected by the chain-breaking operation.

3.2 Markables

Noun-Noun compounding In Universal Dependencies (UD), a noun compound is a relation that connects common noun modifiers. GUM marks the noun compounding span in coreference annotations while OntoNotes annotation guidelines do not permit any common noun compound modifiers. As in (5), the two compounds are marked in one coreference chain.


(5) *Before:* Allergan Inc. said it received approval to sell the PhacoFlex intraocular lens, the first foldable silicone lens available for cataract surgery . The lens' foldability enables it to be inserted in smaller incisions than are now possible for cataract surgery .

(6) *After:* ... available for cataract surgery ... possible for cataract surgery .

To convert the dataset, the conversion program recursively removes the compounding construction from a coreference chain and create a singleton span for that compound, as the two dashed markable spans in (6).

Bridging and split antecedents Bridging coreference occurs when two entities do not corefer exactly, but the basis for the identifiability of one referent is the previous mention of one or more previous referents. Particularly, in GUM, an anaphor may have multiple antecedents and each antecedent creates a part-whole relationship with the referent (split antecedent). This part-whole relation, however, is not a valid coreference relation in OntoNotes because two nominal mentions are not semantically identical.

(7) *Before:* Claire Bailey-Ross ... Andrew Beresford ... Daniel Smith ... In this paper, we report upon the novel insights ... We will discuss the potential implications ...




(8) *After:* Claire Bailey-Ross ... Andrew Beresford ... Daniel Smith ... In this paper, we report upon the novel insights ... We will discuss the potential implications ...

As in the example (7), the three proper nouns link to the collective nominal *we*. With the rich coreference annotation in GUM, all coreference relations identified as BRIDGE, including both split antecedent and other types of bridging anaphora, are removed, possibly leaving the nominals with the original bridging relations as singletons. The example (8) shows the coreference relations after the conversion process. The proper names are unlinked (dashed boxes) from the original cluster while the pronouns are not affected.

Copula Copula predicates are annotated in GUM while are not markables in OntoNotes. If two markables are coreferred, the conversion process utilizes the UD annotations to identify whether or not a copula construction is between the two mention spans. If the ROOT is within the second span and it is the head of COP, we decide to remove the second span from the coreference chain.

(9) *Before:* The viewing experience of art is a complex one ... The time it takes ...



(10) *After:* The viewing experience of art is a complex one ... The time it takes ...

For example, in (9), the head noun *one* in the second span *a complex one* is the ROOT and the head of the copula *is*. Additionally, it corefers with a markable that is the subject of the copula construction. Therefore, we unlink the second span from the cluster and connect the first span with the markable that the second span refers to. The post-conversion annotations are marked in (10).

Nested entity A proper noun may be included in a larger proper noun mention span, as in the nested proper noun *America* within the span *Bank of America*. In OntoNotes, all proper names are considered to be atomic, so that *America* will not be annotated

in the above example. Differently from OntoNotes, GUM allows all nested entities to be considered as valid referents. Because a nominal mention is a candidate referent, its possessive modifier, which is also considered as a candidate referent, should be removed from the coreference relation.

(11) *Before:* I 'm about to go visit and would like to know the best way to communicate with her if it 's helpful .

(12) *After:* I 'm about to go visit and would like to know the best way to communicate with her if it 's helpful .

For example, in (11), *it* and *the best way to communicate with her if it's helpful* refer to the same entity, while the pronoun is within the span of the noun phrase, so we remove *it* as in (12). In general, the conversion process removes the nested entity when it is included by its antecedent.

Singletons Singletons are markables that are not referred to by other mentions in a document. GUM explicitly annotates all nominal spans while OntoNotes only considers co-referring noun phrases as valid mention spans. Exemplified in (13), the dashed box indicates that the span is not referred to in the context.

(13) *Before:* a unique collection of 17th Century Zurbarán paintings

The singleton removal process is the last layer for the conversion because previous steps may delete coreference relations and create new markables that are not referred to again.

3.3 Agreement study

To evaluate conversion accuracy, three annotators, including an original OntoNotes project member, conducted an agreement study on 3 documents, containing 2,500 tokens and 371 output mentions. Re-annotating from scratch based on OntoNotes guidelines, the conversion achieves a span detection score of ~96 and CoNLL coreference score of ~92, approximately the same as human agreement scores on OntoNotes. After adjudication, the conversion was found to make only 8/371 errors, in addition to 2 errors due to mistakes in the original GUM data, meaning that degradation due to conversion errors is marginal, and consistency should be close to the variability in OntoNotes itself.

Genre	MUC			B ³			CEAF _{φ4}			Avg. F1	Mention Detection		
	P	R	F1	P	R	F1	P	R	F1		P	R	F1
	dcoref												
<i>academic</i>	35.1	37.5	36.2	32.6	34.4	33.5	35.7	37.5	36.6	35.4	48.3	51.3	49.8
<i>bio</i>	58.0	61.6	59.8	36.8	43.6	39.9	32.1	33.5	32.8	44.1	58.9	62.3	60.6
<i>conversation</i>	62.2	52.9	57.1	40.5	36.7	38.5	37.1	38.2	37.6	44.4	76.6	67.8	72.0
<i>fiction</i>	57.7	43.9	49.9	50.4	33.2	40.0	37.1	49.0	42.2	44.0	68.2	59.0	63.3
<i>interview</i>	57.3	53.3	55.2	<u>29.3</u>	<u>21.6</u>	<u>24.8</u>	<u>22.4</u>	<u>24.6</u>	<u>23.5</u>	<u>27.6</u>	64.3	60.3	62.2
<i>news</i>	57.6	55.2	56.4	45.7	42.3	44.0	39.6	32.5	35.7	45.3	44.0	50.2	46.9
<i>reddit</i>	59.6	65.1	62.3	38.3	53.5	44.6	32.9	34.0	33.5	46.8	60.5	64.6	62.5
<i>speech</i>	50.6	56.2	53.2	40.1	43.9	41.9	46.5	38.6	42.2	45.8	63.5	64.2	63.9
<i>textbook</i>	36.0	34.2	35.1	32.7	31.0	31.9	23.9	39.9	29.9	32.3	18.1	45.8	26.0
<i>vlog</i>	63.6	69.4	66.4	56.4	60.8	58.5	31.4	36.2	33.6	52.8	76.4	76.8	76.6
<i>voyage</i>	<u>34.7</u>	37.1	35.9	30.7	28.7	29.7	29.7	35.8	32.5	32.7	46.6	62.4	53.3
<i>whow</i>	35.8	<u>24.2</u>	<u>28.9</u>	30.0	24.5	27.0	29.9	34.0	31.8	29.2	50.0	42.9	46.2
All OntoGUM	45.7	47.0	46.3	17.1	38.1	37.6	33.4	37.3	35.3	39.7	56.2	59.1	57.6
OntoNotes	57.5	61.8	59.6	68.2	68.4	68.3	47.7	43.4	45.5	57.8	66.8	75.1	70.7
	Joshi et al. (2019a)												
<i>academic</i>	84.5	53.0	65.1	83.3	48.5	61.3	83.2	47.0	60.1	62.2	91.0	55.2	68.7
<i>bio</i>	85.8	74.7	79.8	61.4	64.3	62.8	65.4	49.9	56.6	66.4	87.7	74.5	80.5
<i>conversation</i>	85.0	73.4	78.7	67.9	64.5	66.2	70.2	51.1	59.1	68.0	93.0	77.9	84.8
<i>fiction</i>	87.0	62.5	73.0	78.8	54.1	64.1	62.5	53.1	57.4	64.8	91.1	67.7	77.7
<i>interview</i>	83.9	71.8	77.4	76.1	60.4	67.3	72.9	50.6	59.7	68.2	85.9	70.4	77.3
<i>news</i>	65.3	65.8	65.5	60.1	59.6	59.9	58.9	54.3	56.5	60.6	71.9	70.5	71.2
<i>reddit</i>	76.7	67.4	71.7	67.5	60.3	63.7	69.5	<u>40.5</u>	<u>51.1</u>	61.7	85.3	68.1	75.8
<i>speech</i>	83.3	63.4	72.0	71.2	56.6	63.1	77.3	57.3	65.8	67.0	91.9	69.4	79.0
<i>textbook</i>	<u>50.0</u>	66.6	57.1	<u>45.2</u>	65.7	53.6	<u>55.6</u>	55.6	55.6	<u>55.5</u>	60.0	72.2	65.5
<i>vlog</i>	86.1	86.1	86.1	78.4	79.8	79.1	63.6	47.7	54.5	73.3	89.4	85.4	87.4
<i>voyage</i>	69.0	70.4	69.7	52.7	64.1	57.9	65.9	53.0	58.8	62.1	78.9	75.5	77.2
<i>whow</i>	84.8	<u>40.9</u>	<u>55.2</u>	83.4	<u>39.2</u>	<u>53.3</u>	71.4	57.4	63.6	57.4	93.2	52.4	67.1
All OntoGUM	79.7	66.3	72.4	69.5	58.58	63.7	67.7	50.7	58.0	64.6	85.4	69.2	76.5
OntoNotes	85.8	84.8	85.3	78.3	77.9	78.1	76.4	74.2	75.3	79.6	89.1	86.5	87.8

Table 3: Results on the OntoGUM’s test dataset with the deterministic coref model (top) and the SOTA coreference model (bottom). The underlined text is the lowest score across 12 genres and bold text is the highest.

4 Experiments

We evaluate two systems on the 12 OntoGUM genres, using the official CoNLL-2012 scorer (Pradhan et al., 2012, 2014). The primary score is the average F1 of three metrics – MUC, B³, and CEAF_{φ4}.

Deterministic coreference model We first run the deterministic system (dcoref, part of Stanford CoreNLP, Manning et al. 2014) on the OntoGUM benchmark, as it remains a popular option for off-the-shelf coreference resolution. As a rule-based system, it does not require training data, so we directly test it on OntoGUM’s test set. However, POS tags, lemmas, and named-entity (NER) information are predicted by CoreNLP, which does have a domain bias favoring newswire. The system’s multi-sieve structure and token-level features such as gender and number remain unchanged. We expect that the linguistic rules will function similarly across datasets and genres, notwithstanding biases of the tools providing input features to those rules.

SOTA neural model Combining the e2e approach with a contextualized LM and span masking is the current SOTA on OntoNotes. The system utilizes the pretrained SpanBERT-large model, fine-tuned on the OntoNotes training set. Hyperparameters are identical to the evaluation of OntoNotes test to ensure comparable results between the benchmarks. We note that while we choose the SOTA system as a ‘best case scenario’, most off-the-shelf neural NLP toolkits (e.g. spaCy) actually use somewhat simpler e2e models than SpanBERT-large, due to memory/performance constraints.

5 Results

OntoGUM vs. OntoNotes The last rows in each half of Table 3 give overall results for the systems on each benchmark. e2e+SpanBERT encounters a substantial degradation of 15 points (19%) on OntoGUM, likely due to lower test set lexical and stylistic overlap, including novel mention pairs. We note that its average score of 64.6 is somewhat optimistic, especially given that the system receives access to gold speaker information wherever avail-

Genres	PRON (R)	Other (R)	Total	CoNLL	Span
<i>vlog</i>	600 (.66)	309 (.34)	909	1	1
<i>interview</i>	1223 (.45)	1485 (.55)	2708	2	6
<i>conversation</i>	729 (.61)	323 (.39)	1052	3	2
<i>speech</i>	245 (.40)	364 (.60)	609	4	4
<i>bio</i>	796 (.34)	1529 (.66)	2325	5	3
<i>fiction</i>	1700 (.61)	1091 (.39)	2791	6	5
<i>academic</i>	262 (.21)	997 (.79)	1259	7	10
<i>voyage</i>	300 (.22)	1053 (.78)	1353	8	7
<i>reddit</i>	1337 (.55)	1077 (.45)	2414	9	8
<i>news</i>	340 (.19)	1483 (.81)	1823	10	9
<i>whow</i>	1001 (.47)	1129 (.53)	2130	11	11
<i>textbook</i>	165 (.34)	315 (.66)	480	12	12

Table 4: Mention-type counts (ratios) & ranks of SOTA scores by genre (CoNLL score + span detection).

able (including in *fiction*, *conversation* and *interview*, some of the better scoring genres), which is usually unrealistic. *dcoref*, assumed to be more stable across genres, also sees losses on OntoGUM of over 18 points (30%). We believe at least part of the degradation may be due to mention detection, which is trained on different domains for both systems (see the last three columns in the table). These results suggest that input data from CoreNLP degrades substantially on OntoGUM, or that some types of coreferent expressions in OntoGUM are linguistically distinct from those in OntoNotes, or both, making OntoGUM a challenging benchmark for systems developed using OntoNotes.

Comparing genres Both systems degrade more on specific genres. For example, while *vlog* (with gold speaker information) fares well for both systems, neither does well on *textbook*, and even the SOTA system falls well below (or around) 60s for the *news*, *whow* and *textbook* genres. This might be surprising for *vlog*, which contains transcripts of spontaneous unedited speech from YouTube Creative Commons vlogs quite unlike OntoNotes data; conversely the result is less expected for carefully edited texts which are somewhat similar to data in OntoNotes: OntoNotes contains roughly 30% newswire text, and it is not immediately clear that GUM’s *news* section, which comes from recent Wikinews articles, differs much in genre. Examples (14)–(15) illustrate incorrectly predicted coreference chains from both sources and the type of language they contain.

- (14) *I’ve been here just crushing ultrasounds ... I’ve been like crushing these all day today ... I got sick when I was on Croatia for vacation. I have no idea what it says, but I think they’re cough drops.* (example from a radiologist’s vlog, incorrect: ultrasounds

≠ cough drops)

- (15) The report has prompted calls for all edible salt to be iodised ... Tasmania was excluded from the study - where a voluntary iodine fortification program using iodised salt in bread, is ongoing (newswire example, incorrect span and coref: [the study - where a voluntary...])

These examples show that errors occur readily even in quite characteristic news writing, while genre disparity by itself does not guarantee low performance, as in the case of the vlogs whose language is markedly different. In sum, these observations suggest that accurate coreference for downstream applications cannot be expected in some common well edited genres, despite the prevalence of news data in OntoNotes (albeit specifically from the Wall Street Journal, around 1990). This motivates the use of OntoGUM as a test set for future benchmarking, in order to give the NLP community a realistic idea of the range of performance we may see on contemporary data ‘in the wild’.

We also suspect that prevalence of pronouns and gold speaker information produce better scores in the results. Table 4 ranks genres by their e2e CoNLL score, and gives the proportions of pronouns, as well as score rankings for span detection. Because pronouns are usually easier to detect and pair than nouns (Durrett and Klein, 2013), more pronouns usually means higher scores. On genres with more than 50% pronouns and gold speakers (*vlog*, *interview*, *conversation*, *speech*, *fiction*) e2e gets much higher results, while genres with few pronouns (<30%) have lower scores (*academic*, *voyage*, *news*). This diversity over 12 genres supports the usefulness of OntoGUM, which can evaluate the generalizability of coreference systems.

6 Conclusion

This paper presented the mechanics of the conversion of the GUM corpus to the OntoGUM version using the OntoNotes coreference scheme, creating the largest open, gold standard coreference dataset following the OntoNotes scheme, which adds several new genres (including more spoken data) to the OntoNotes family. The corpus is automatically converted from GUM by modifying the existing markable spans and coreference relations using multi-layer annotations, such as dependency trees. Results showed a lack of generalizability of existing systems, especially in genres low in pronouns and lacking speaker information. We suspect that at least part of the success of SOTA approaches is due to correct mention detection and high matching scores in genres rich in pronouns, and more so with gold speaker information. Success for other types of mentions in OntoNotes data appears to be much more sensitive to lexical features, performing well on the benchmark test set with high lexical overlap to the training data, but degrading very substantially outside of it, even on newswire texts from our OntoGUM data. This supports use of this challenging dataset for future work, which we hope will benefit evaluations of systems targeting the OntoNotes standard.

References

- Hong Chen, Zhenhua Fan, Hao Lu, Alan Yuille, and Shu Rong. 2018. [PreCo: A large-scale dataset in preschool vocabulary for coreference resolution](#). In *Proceedings of EMNLP 2018*, pages 172–181, Brussels.
- Kevin Clark and Christopher D. Manning. 2015. [Entity-centric coreference resolution with model stacking](#). In *Proceedings of ACL-IJCNLP 2015, Long Papers*, pages 1405–1415, Beijing, China.
- Kevin Clark and Christopher D. Manning. 2016a. [Deep reinforcement learning for mention-ranking coreference models](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2256–2262, Austin, Texas.
- Kevin Clark and Christopher D. Manning. 2016b. [Improving coreference resolution by learning entity-level distributed representations](#). In *Proceedings of ACL 2016, Long Papers*, pages 643–653, Berlin.
- Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. [Universal Dependencies](#). *Computational Linguistics*, 47(2):255–308.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of NAACL 2019*, pages 4171–4186, Minneapolis, MN.
- Greg Durrett and Dan Klein. 2013. [Easy victories and uphill battles in coreference resolution](#). In *Proceedings of EMNLP 2013*, pages 1971–1982, Seattle, WA.
- Greg Durrett and Dan Klein. 2014. [A joint model for entity analysis: Coreference, typing, and linking](#). *TACL*, 2:477–490.
- Hongliang Fei, Xu Li, Dingcheng Li, and Ping Li. 2019. [End-to-end deep reinforcement learning based coreference resolution](#). In *Proceedings of ACL 2019*, pages 660–665, Florence, Italy.
- Abbas Ghaddar and Phillippe Langlais. 2016. [Wiki-Coref: An English coreference-annotated corpus of Wikipedia articles](#). In *Proceedings of LREC 2016*, pages 136–142, Portorož, Slovenia.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2019a. [SpanBERT: Improving pre-training by representing and predicting spans](#). *CoRR*, abs/1907.10529.
- Mandar Joshi, Omer Levy, Daniel S. Weld, and Luke Zettlemoyer. 2019b. [BERT for coreference resolution: Baselines and analysis](#). In *Proceedings of EMNLP-IJCNLP 2019*, pages 5803–5808, Hong Kong, China.
- 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 EMNLP 2017*, pages 188–197, Copenhagen, Denmark.
- Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. [Higher-order coreference resolution with coarse-to-fine inference](#). In *Proceedings of NAACL 2018, Short Papers*, pages 687–692, New Orleans, LA.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *ACL 2014 System Demonstrations*, pages 55–60.
- Nafise Sadat Moosavi and Michael Strube. 2017. [Lexical features in coreference resolution: To be used with caution](#). In *Proceedings of ACL 2017, Short Papers*, pages 14–19, Vancouver, Canada.
- Nafise Sadat Moosavi and Michael Strube. 2018. [Using linguistic features to improve the generalization capability of neural coreference resolvers](#). In *Proceedings of EMNLP 2018*, pages 193–203, Brussels.

- Massimo Poesio, Yulia Grishina, Varada Kolhatkar, Nafise Moosavi, Ina Roesiger, Adam Roussel, Fabian Simonjetz, Alexandra Uma, Olga Uryupina, Juntao Yu, and Heike Zinsmeister. 2018. [Anaphora resolution with the ARRAU corpus](#). In *Proceedings of CRAC 2018*, pages 11–22, New Orleans, LA.
- Massimo Poesio, Roland Stuckardt, and Yannick Versley. 2016. *Anaphora resolution*. Springer.
- Sameer Pradhan, Xiaoqiang Luo, Marta Recasens, Eduard Hovy, Vincent Ng, and Michael Strube. 2014. [Scoring coreference partitions of predicted mentions: A reference implementation](#). In *Proceedings of ACL 2014*, pages 30–35, Baltimore, MD.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. [Towards robust linguistic analysis using OntoNotes](#). In *Proceedings of CoNLL 2013*, pages 143–152, Sofia.
- 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.
- Marta Recasens, Marie-Catherine de Marneffe, and Christopher Potts. 2013. [The life and death of discourse entities: Identifying singleton mentions](#). In *Proceedings of NAACL 2013*, pages 627–633, Atlanta, Georgia.
- Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. 2018. [Mind the GAP: A balanced corpus of gendered ambiguous pronouns](#). *TACL*, pages 605–617.
- Ralph Weischedel, Eduard Hovy, Mitchell Marcus, Martha Palmer, Robert Belvin, Sameer Pradhan, Lance Ramshaw, and Nianwen Xue. 2011. [OntoNotes: A large training corpus for enhanced processing](#). In Joseph Olive, Caitlin Christianson, and John McCary, editors, *Handbook of Natural Language Processing and Machine Translation: DARPA Global Autonomous Language Exploitation*. Springer.
- Sam Wiseman, Alexander M. Rush, Stuart Shieber, and Jason Weston. 2015. [Learning anaphoricity and antecedent ranking features for coreference resolution](#). In *Proceedings of ACL-IJCNLP 2015, Long Papers*, pages 1416–1426, Beijing, China.
- Sam Wiseman, Alexander M. Rush, and Stuart M. Shieber. 2016. [Learning global features for coreference resolution](#). In *Proceedings of NAACL 2016*, pages 994–1004, San Diego, CA.
- 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. 2021. [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.