

# A Dataset of Literary Coreference

David Bamman, Olivia Lewke and Anya Mansoor

School of Information

University of California, Berkeley

{dbamman, olivialewke, anyamansoor}@berkeley.edu

## Abstract

We present in this work a new dataset of coreference annotations for works of literature in English, covering 29,103 mentions in 210,532 tokens from 100 works of fiction. This dataset differs from previous coreference datasets in containing documents whose average length (2,105.3 words) is four times longer than other benchmark datasets (463.7 for OntoNotes), and contains examples of difficult coreference problems common in literature. This dataset allows for an evaluation of cross-domain performance for the task of coreference resolution, and analysis into the characteristics of long-distance within-document coreference.

**Keywords:** coreference, corpus annotation, literature

## 1. Introduction

Modern coreference resolution systems are typically evaluated on the benchmark OntoNotes dataset (Hovy et al., 2006) originally used for the shared tasks on modeling unrestricted coreference at CoNLL 2011 (Pradhan et al., 2011) and CoNLL 2012 (Pradhan et al., 2012). Nearly all modern systems evaluate on this data exclusively (Lee et al., 2017; Lee et al., 2018; Wiseman et al., 2016; Clark and Manning, 2016; Kantor and Globerson, 2019).

However, the domains included within OntoNotes are relatively narrow—predominantly focused on news (broadcast, magazine, and newswire), conversation, the Bible, and the web—and much recent work has shown the inability of coreference resolution systems to generalize to new domains beyond those they were trained on (Moosavi and Strube, 2017; Moosavi and Strube, 2018; Subramanian and Roth, 2019; Lestari, 2018). While coreference datasets exist for other domains such as Wikipedia (Ghaddar and Langlais, 2016), scientific articles (Schäfer et al., 2012) and school examinations (Chen et al., 2018), one area lacking robust data is English literature. Since a growing body of research leveraging computational methods to reason about literature is increasingly relying on coreference information—for example, in the study of literary character (Underwood et al., 2018; Kraicer and Piper, 2018)—an annotated dataset in this domain is critical for robustly assessing the quality of systems used.

Additionally, as Roesiger et al. (2018) have outlined, literary texts exhibit markedly different behavior with respect to coreference, rooted in the differing narrative spheres between the world of the narrator and that of the entities in the narration; the level of alternation between generic and specific mentions; the evolution of a character over the space of a long novel; and the differing levels of knowledge that the characters and readers have of the factuality of the events being described. Each of these phenomena has the potential to impact coreference performance if modeled with a system trained on data where it is rarely seen.

We present in this work a new dataset of coreference annotations for literary texts, to help assess the performance of coreference resolution systems for this domain. We make the following contributions:

- We present a new dataset of 210,532 tokens from 100 different literary novels, annotated for coreference relations between the ACE categories of people, locations, organizations, facilities, geopolitical entities and vehicles (Walker et al., 2006).
- We characterize the behavior of coreference phenomena within this dataset, where the average document length is 2,000 words. This long document length allows us to investigate the temporal distance over which entities exist in discourse, model the burstiness of major entities, and measure the distribution in distance to closest antecedents for different categories of entities; all of these measures can help inform future coreference resolution systems in this domain.
- We assess the performance of a neural coreference resolution system on this data while varying the domain of the training source. While a model trained on in-domain literary data achieves an F-score six points higher than one trained on OntoNotes, it performs equivalently to a model trained on PreCo (Chen et al., 2018)—a very different domain, but two orders of magnitude larger.

## 2. Related work

While OntoNotes (Hovy et al., 2006) has been the *de facto* benchmark dataset used for coreference resolution, several other resources have been created to explore coreference in other domains—including Wikipedia articles in WikiCoref (Ghaddar and Langlais, 2016); articles in the ACL Anthology (Schäfer et al., 2012); school examinations in PreCo (Chen et al., 2018); Quiz Bowl questions (Guha et al., 2015); and several different aspects of the biomedical domain (Cohen et al., 2017; Nguyen et al., 2011; D’Souza and Ng, 2012; Yang et al., 2005; Gasperin et al., 2007; Su et al., 2008).

Less work attends to coreference within specifically literary texts in English, both in terms of annotated data and models for coreference resolution in literature. Bamman et al. (2014) annotates pronominal coreference links within a small sample of five novels for coreference resolution within BookNLP; Vala et al. (2016) annotates coreference in *Pride*

and *Prejudice*, while Vala et al. (2015) develops a method leveraging character and alias information for character detection in *Sherlock Holmes*, *Pride and Prejudice* and *The Moonstone*. The GUM corpus (Zeldes, 2017) contains approximately 12,000 tokens of fiction in its multilayer annotations (which include coreference). Roesiger et al. (2018) describes a general framework for conceptualizing coreference in literary texts, and presents a tool for annotating coreference in long documents. Comparatively larger resources, however, exist for German literature—Krug et al. (2017) create a dataset of coreference annotations between character mentions in German novels, which helped assess a rule-based system for coreference resolution (Krug et al., 2015).

At the same time, coreference in being used as the foundation for empirical arguments in the humanities, including studies that leverage coreference information in order to identify characters and take measurements about them—such as the amount of attention given to characters as a function of their gender (Underwood et al., 2018; Kraicer and Piper, 2018). Several studies have used coreference and alias resolution in the service of identifying character networks (Elson et al., 2010; Agarwal et al., 2012; Lee and Yeung, 2012; Sudhahar and Cristianini, 2013; Jannidis et al., 2016; Piper et al., 2017).

### 3. Data

We draw our source material for annotation from the texts in LitBank (Bamman et al., 2019), which consists of 210,532 tokens drawn from 100 different works of English-language fiction from 1719 to 1922 (all within the US public domain). These texts include a range of literary styles, and include existing annotations for ACE-style entities (people, organizations, locations, geo-political entities, facilities and vehicles), along with *realis* events (Sims et al., 2019). One motivating factor in the corpus selection process was to select a significant sample (approximately 2,000 words) across a wide range of different texts in order to test the performance of systems across different works and authors within the broader domain of literature.

### 4. Annotation

Our annotation style largely follows that of OntoNotes, in defining the boundaries for markable mentions that can be involved in coreference and in defining the criteria for establishing coreference between them. However, we make several important departures from their guidelines to accommodate the specific phenomena encountered in literary texts, noted here.

#### 4.1. Markables

We make several departures from the OntoNotes definition of a markable span—one that is eligible to be involved in coreference.

##### 4.1.1. Singletons

OntoNotes does not allow singleton mentions (noun phrases that are not involved in coreference) to be markable; this decision complicates coreference resolution in test documents, where a separate preprocessing step must be carried out to

define the set of candidate noun phrases that are eligible for coreference (i.e., to decide if a noun phrase is a singleton or part a coreference chain). In the context of OntoNotes, this decision is justified by the presence of a full syntactic parse for all documents; every NP non-terminal in the parse tree that is not marked as participating in a coreference chain can be inferred to be a singleton. In our work, we **do** treat all singleton mentions as markable, removing the need for a gold syntactic parse for each document. 17.4% of all mentions are singletons in our data, in contrast to the 56% observed by Recasens et al. (2013) in OntoNotes.

##### 4.1.2. Entity types

While OntoNotes covers unrestricted coreference, we limit the markable entities to only those categories annotated in LitBank. The existing annotations in Litbank cover six entity types: people (PER), facilities (FAC), locations (LOC), geo-political entities (GPE), organizations (ORG), and vehicles (VEH). As table 1 shows, the majority (83.1%) of mentions among these annotations are people, attesting to the significant focus on characters in this domain.

Category	<i>n</i>	Frequency
PER	24,180	83.1%
FAC	2,330	8.0%
LOC	1,289	4.4%
GPE	948	3.3%
VEH	207	0.7%
ORG	149	0.5%

Table 1: Counts of entity type.

##### 4.1.3. Entity categories

Like OntoNotes, we include noun phrases that are proper names (PROP), common phrases (NOM), and pronouns (PRON). The entity annotations in LitBank, however, were originally designed for the task of entity tagging, and focus only on proper nouns (*Tom Sawyer*) and common nouns (*the boy*). To correct for this, we include as markable spans all personal pronouns as well, including common forms (*I, me, my, myself, you, your, yourself, she, her, herself, he, him, his, himself, it, its, we, our, they, them, their*), historical forms (*thou, thee, thine, ye*) and forms originating in transcriptions of speech (*'em, 'ee, yeh, yer*). While annotating a mention, we record its entity category (PROP, NOM, PRON) to enable more fine-grained analysis. As table 2 illustrates, pronouns account for the majority of mentions in literary texts.

Category	<i>n</i>	Frequency
PRON	15,816	54.3%
NOM	9,737	33.5%
PROP	3,550	12.2%

Table 2: Counts of entity category.

##### 4.1.4. Quantified and negated noun phrases

In order to capture potential coreference chains such as “[No mother]<sub>x</sub> should be separated from [her]<sub>x</sub> child”, we treat as markables all quantified and negated noun phrases, including those modified by *some, any, many, few, no, neither*, and

*none of*. While OntoNotes rarely annotates these since they would result in singleton coreference chains, we annotate all examples in our dataset.

#### 4.1.5. Maximal spans

Following OntoNotes, we mark the maximal extent of a span, as in the following:

[The boy who painted the fence and ate lunch] ran away.

This is notably distinct from the work of Krug et al. (2017), who only annotate the syntactic heads of characters, and not their full extent.

#### 4.1.6. Honorifics

While ACE and LitBank both annotate honorifics in personal names (such as “[Mr.] Collins”, and “[Miss] Havisham]”), these are often stripped by entity recognition systems and coreference resolution systems. We exclude those from being distinct markable spans here, leaving as entities only the entire maximal span “[Mr. Collins]” and “[Miss Havisham]”.

## 4.2. Coreference

As Roesiger et al. (2018) point out, there are number of ways in which the style of literary texts influences coreference in a manner different from the domains captured in existing datasets. We outline several of those distinctions, along with the decisions we make in annotating them.

### 4.2.1. Generic vs. specific mentions

Most existing datasets for coreference make a distinction between *generic* and *specific* mentions of entities. A generic noun phrase refers to “a kind or class of individuals” (Reiter and Frank, 2010), while a specific mention refers to a specific individual or group. OntoNotes allows generic mentions to be coreferent only with pronominals and not other generic mentions (“[Doctors]<sub>x</sub> care for [their]<sub>x</sub> patients); PreCo additionally allows generic mentions to be coreferent directly (“[Doctors]<sub>x</sub> do what is best for [doctors]<sub>x</sub>”).

We follow the OntoNotes decision here, and disallow generic mentions from being coreferent with each other; importantly, this impacts the annotation of generic “you”, either used impersonally:

Right and left, the streets take [you]<sub>x</sub> waterward  
(Melville, *Moby Dick*)

Or in an address to the reader:

Let me now take [you]<sub>x</sub> on to the day of the assault  
(Collins, *The Moonstone*).

One area, however, where the generic/specific distinction becomes difficult is in the use of a nominally generic phrase that may refer ambiguously to a specific individual, as in the following example:

Whereas with respect to Turkey, I had much ado to keep him from being a reproach to me. His clothes were apt to look oily and smell of eating-houses. He wore his pantaloons very loose and

baggy in summer. His coats were execrable; his hat not to be handled. But while the hat was a thing of indifference to me, inasmuch as his natural civility and deference, as a dependent Englishman, always led him to doff it the moment he entered the room, yet his coat was another matter. Concerning his coats, I reasoned with him; but with no effect. The truth was, I suppose, that [a man of so small an income] could not afford to sport such a lustrous face and a lustrous coat at one and the same time. (Melville, *Bartleby, The Scrivener*)

Here Melville describes the physical appearance of the character Turkey in *Bartleby, The Scrivener*; while the generic entity “a man of so small an income” surely is being used to describe Turkey, it refers to a more general class of entity, of which Turkey is only a member.

We can see this to be an example of *class near-identity* (Recasens et al., 2010), where “[Turkey]” and “[a man of so small an income]” share an *is-a* relation, with latter being a more general realization of the former. In these cases, while the generic mention may be interpreted as referring to a specific individual, we still treat it as subject to the guidelines for generics, and disallow coreference between the generic mention and any specific individual.

### 4.2.2. Copulae


Existing datasets like OntoNotes and PreCo differ in their annotation of copulae—structures where a given entity is linked to an attribute through a copular verb such as *be*, *appear*, *feel*, *seem*, etc, as in “[John] is [a doctor]”. Datasets either rule out their eligibility for coreference at all (as in OntoNotes) or as being fully coreferent with an entire coreference chain (PreCo). While PreCo allows copular attributes to be coreferent with their subject, we follow OntoNotes in seeing the attributes in copular structures as not fundamentally referential in nature; it is the act of predication that associates two mentions, not their referentiality. We can see this in part by examining the range of copular structures present in our literary dataset, in which we can clearly see that the strength of the association between a subject and attribute is moderated by the form of the linking verb (and any adverbials that modify it). While example 1 below is a copular structure that clearly asserts the equivalence of the comparands, 2 is not—“in many ways” suggests that the predication is not complete (there exist some ways in which he was not “a most valuable person”). Examples 3 and 4 are copular structures about assertions in the past, which may no longer hold in the present; and examples 5 and 6 are both copular but negated.

1. [I] am [a rather elderly man] (Melville, *Bartleby*).
2. [He] was in many ways [a most valuable person to me]. (Melville, *Bartleby*).
3. Mrs. Marjoribanks, poor lady, had been an invalid for many years (Oliphant, *Miss Marjoribanks*).
4. [It] was [a convent] before the Thirty Year’s War (Von Arnim, *Elizabeth and Her German Garden*).

5. [She] was *not* to be described as [a tall girl] (Oliphant, *Miss Marjoribanks*).
6. [Lucilla] was *not* [the woman to be disconcerted] (Oliphant, *Miss Marjoribanks*).


The ways in which a copula can serve as an identity-linking function between two mentions is clearly on a continuum. To capture just those that are identity-establishing, we only annotate examples like 1 above: only those that are asserted as currently holding true, not relations that might hold in the future, hypotheticals, or negation.

And since it is the copular structure itself that asserts identity rather than the referentiality of the mention, we annotate copulae as a link between individual mentions (as in the OntoNotes treatment of appositions), so that example (1) would be annotated as the following:



#### 4.2.3. Apposition

Like OntoNotes, we also annotate apposition as a distinct relation that holds between specific mentions that are immediately adjacent:



#### 4.2.4. Identity and Near-Identity

Classical models of coreference often characterize the problem as deciding whether two mentions refer to the same entity in the real world, which can easily become entangled in metaphysical complexities on the nature of identity (for an overview, see Gallois (2016)). While all discourse involves such complexities, it is exacerbated in literary novels, which not only describe entities that may exist, and change, over the course of a long narrative timeline, but may also describe the historical background for those entities. Novels may describe a city as it evolves over the course of a millennium (such as London), or the development of a child into an adult over a period of decades (*Great Expectations*). Does the London of 1922 refer to the same entity as London in 1599? And does the seven-year-old Pip at the beginning of *Great Expectations* refer to the same entity as the thirty-year-old Pip at the end?

Rather than determine the identity of reference of real-world entities, we draw on the formalization of Recasens et al. (2011) in their discussion of near-identity in coreference, in which they outline the specific operations of neutralization/compression and refocusing/decompression that respectively minimize or maximize the differences in feature values between two discourse entities, effectively making them *more* identical in the case of neutralization and *less* identical in the case of refocusing. Importantly, the degree of identity in coreference here is tied not to entities that exist in the real world, but rather to discourse entities constructed in a discourse, and for whom a specific pragmatic context may encourage two mentions to be seen as more

or less coherent—independent of any real-world status they might have (Nunberg, 1984). Recasens et al. (2011) cite the example of *Postville* to illustrate pragmatic near-identity:

On homecoming night [Postville]<sub>x</sub> feels like Hometown, USA, but a look around [this town of 2,000]<sub>x</sub> shows it's become a miniature Ellis Island. This was an all-white, all-Christian community . . . For those who prefer [the old Postville]<sub>y</sub>, Mayor John Hyman has a simple answer. (Recasens et al., 2011, 10)

Here *Postville* and *the old Postville* are pragmatically distanced through an act of refocusing to emphasize their differences (the inhabitants of old Postville being predominantly white, while the Postville contemporaneous with writing is more diverse).

In Hawthorne's *House of the Seven Gables* we see a similar case of near-identity involving entities that exist with the same physical boundaries but different temporal extent; here, however, we see a reverse effect of compression, flattening the temporal differences between the eponymous *House of the Seven Gables* and *the old Pyncheon House*:

Halfway down a by-street of one of our New England towns stands [a rusty wooden house, with seven acutely peaked gables, facing towards various points of the compass, and a huge, clustered chimney in the midst]<sub>x</sub>. The street is Pyncheon Street; [the house]<sub>x</sub> is [the old Pyncheon House]<sub>cop→the house</sub>.

Here, even though the houses have different feature values (different names and different temporal extents), compression equates them through the use of copular predication. At the same time, this entity is distanced through refocusing with respect to an temporally earlier entity that occupied the same physical location:

[The House of the Seven Gables]<sub>x</sub>, antique as it now looks, was not [the first habitation erected by civilized man on precisely the same spot of ground]. . . . A natural spring of soft and pleasant water . . . had early induced Matthew Maule to build [a hut, shaggy with thatch, at this point, although somewhat too remote from what was then the centre of the village]<sub>y</sub>.

We can see another example of this compression in Orczy's *The Scarlet Pimpernel*:

During the greater part of the day the guillotine had been kept busy at its ghastly work: all that [France]<sub>x</sub> had boasted of in the past centuries, of ancient names, and blue blood, had paid toll to [her]<sub>x</sub> desire for liberty and for fraternity.

*The Scarlet Pimpernel* is set in France during the French Revolution, just after the regime change the revolution entailed. Under a metaphysical criterion of identity, we might distinguish between France as a monarchy prior to revolution and France as a republic under the Committee of

Public Safety (assuming we think of a state as being comprised not only of its territorial holdings, but the structure of its government and the customs/practices of its people). Here, however, Orczy neutralizes those shades of meaning by stressing the France that has persisted through centuries; we treat all mentions of France as coreferent under this compression.

#### 4.2.5. Revelation of identity

One other phenomenon that is more characteristic of literary texts than texts from other domains is a revelation of identity, when one entity is revealed to be identical with another. This revelation can be sudden, in which entities  $x$  and  $y$  were depicted as being separate but are revealed to be the same, as in the case of *whodunnit* detective novels and mysteries. In *Great Expectations*, Pip encounters a convict as a child at the beginning of the novel, and is later supported as a young man by an unknown benefactor; these are revealed to be the same entity at the end of the novel. Here we adopt an approach similar to Roesiger et al. (2018), which annotates identity from the reader’s point of view (as distinct from characters’); here, we can see that the convict and Pip’s benefactor are identical if we presume that a reader’s point of view is scoped over the entirety of the text.

Revelations of identity can also arise through a gradual accumulation of information. In Conan Doyle’s *Adventures of Sherlock Holmes*, Holmes receives a letter noting the imminent arrival of a visitor:

“There will call upon you to-night, at a quarter to eight o’clock,” it said, “[a gentleman who desires to consult you upon a matter of the very deepest moment]<sub>x</sub>. . . . Be in your chamber then at that hour, and do not take it amiss if [your visitor]<sub>x</sub> wear a mask.”

Here the fact that the identity of the visitor described in the letter is the same as that of the author of the letter is not known to the reader; Sherlock then proceeds to deduce that the author of the letter is likely a German man and is the same individual described as the imminent visitor, concluding:

“It only remains, therefore, to discover what is wanted by [this German who writes upon Bohemian paper and prefers wearing a mask to showing [his]<sub>x</sub> face]<sub>x</sub>. And here [he]<sub>x</sub> comes, if I am not mistaken, to resolve all our doubts.”

The equivalence of these entities is pragmatically determined only after several mentions (and realized syntactically through a conjunction). While a reader could not know that the visitor is the same individual as the author when reading the letter above, we treat all mentions as coreferent if a reader can determine that they are identical at any point in the narrative. In other words, we annotate from the perspective of the textual reality, rather than from the reader’s state of knowledge at the moment of encountering the mention to be resolved.

### 4.3. Annotation process

We carry out annotations in a two-step process: first using a custom command-line interface for linking mentions to

entities in text<sup>1</sup> and then transferring the annotations to a BRAT (Stenetorp et al., 2012) GUI interface for checking. All annotations were performed by three annotators (all three authors) after an initial phase of guideline design and consistency checking. While each text was annotated by a single author, we double-annotated a total of 10 full texts (10% of the total collection) in order to assess the consistency of annotation across different annotators. As with prior work, we calculate inter-annotator agreement using the same coreference metrics used for assessing the performance of coreference resolution systems; while OntoNotes reports an average inter-annotator MUC score of 83.0 across its seven subdomains (Pradhan et al., 2012) and PreCo reports a score of 77.5, we find very high agreement between our trained annotators with a MUC score of 95.5.<sup>2</sup> We suspect this is due to a combination of two factors: the smaller pool of high-trained annotators we use for our smaller collection, and the restricted nature of the entities we are labeling (only people, facilities, locations, geopolitical entities, organizations and vehicles) instead of unrestricted coreference. This data is publicly available as a part of LitBank at <https://github.com/dbamman/litbank>.

## 5. Analysis

We carry out several analyses to investigate the behavior of coreference within this dataset.

### 5.1. Spread

Figure 1 examines the distribution in the temporal distance (in narrative time) in which a given entity is active, where distance is measured as the number of tokens between the first mention of an entity and the last mention. The vast majority of entities span relatively short time scales; 50% of them span 173 or fewer tokens (approximately half a page in a printed book). These entities include both generic mentions (*a warrior*), known entities grounded in the real world (*New York*) and specific entities mentioned only briefly that do not factor prominently in the broader discourse.




Figure 1: Spread between first and last mention of entity, in tokens.

In figure 1, each entity is counted equally as a single data point, giving equal weight to major characters and mi-

<sup>1</sup> <https://github.com/dbamman/cl-coref-annotator>

<sup>2</sup> The inter-annotator agreement for all three commonly used coreference metrics are: 93.81  $B^3$ , 95.53 MUC, and 87.72 CEAF <sub>$\phi_4$</sub> .

nor characters. Figure 2 in contrast examines the same phenomenon from the perspective of individual mentions, where each entity is weighted by the number of mentions associated with it. As can be seen, while most entity *types* span relatively short distances, most entity *tokens* are part of coreference chains that span the entire document (ca. 2,000 words). Major characters, in particular, constitute the majority of coreferential mentions.




Figure 2: Spread between first and last mention of entity, in tokens, weighted by the number of mentions in each entity.

## 5.2. Burstiness

To explore whether entities that span long text ranges tend to cluster together in tight bursts, we quantify the burstiness of an entity by measuring the entropy of the empirical distribution defined over mentions in narrative time: we divide the narrative time of a book into 100 equal-sized segments (each spanning roughly 200 words), and define  $d_{e,i}$  to be the relative frequency with which entity  $e$  was mentioned in segment  $i$  (relative to the total number of mentions of  $e$  in the document). We calculate entropy as  $H(d_e)$ .

Figure 3 illustrates this bursty behavior visually by selecting the entities with the lowest entropy (Basil Hallward in *The Picture of Dorian Gray*) and highest entropy (the narrator in *Gulliver's Travels*) among those spanning at least 1500 tokens and mentioned at least 100 times. While Basil Hallward goes through several periods of not being mentioned followed by increased focus (bursty behavior), even the narrator of *Gulliver's Travels* exhibits bursty behavior despite being more uniformly mentioned, with a gap of several hundred words in which he is not mentioned between his otherwise constant focus. Even entities that have high entropy (which should be closer to a uniform distribution of mentions over time) still exhibit bursty behavior in which there is a period of time where they are not mentioned.

## 5.3. Distance to nearest antecedent

Finally, we examine the distribution in distances to the closest antecedent for proper nouns, common nouns and pronouns for non-singleton mentions, as depicted in figure 4. Proper nouns have a median distance of 5 entities to their nearest antecedent, with 90% of antecedents appearing within 42 mentions, and 95% of antecedents within 78.3 mentions. Common nouns (3572 mentions) have a median distance of 6 entities to their nearest antecedent, with 90% of antecedents appearing within 59 mentions, and 95% of




Figure 3: Long-range entities are bursty; the distribution of mentions over narrative time for the entity with the lowest entropy (top; Basil Hallward in Wilde's *The Picture of Dorian Gray*) and highest entropy (bottom; the narrator in Swift's *Gulliver's Travels*).

antecedents within 101 mentions. Pronouns have a median distance of 2 entities to their nearest antecedent, with 90% of antecedents appearing within 5 mentions, and 95% of antecedents within 9 mentions. While coreference systems often impose strict limits to the number of maximum antecedents to consider for long documents (Lee et al., 2017), or use coarse-to-fine inference for reducing the number of candidates (Lee et al., 2018), this suggests that pronouns (which again account for over half of all potentially coreferential mentions in this data) only need to consider a far shorter number of antecedents.

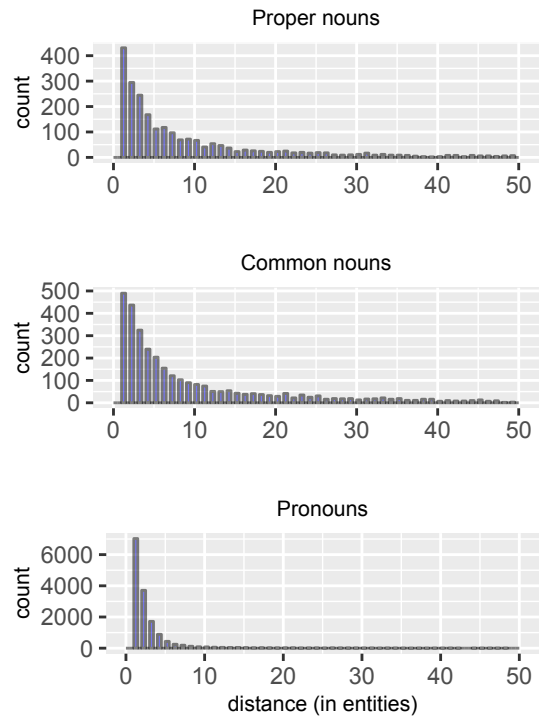


Figure 4: Distance to antecedent in entities.

## 6. Empirical performance

The existence of a dataset of coreference in works of English literature allows us to test the performance of coreference resolution systems in this domain—including those trained on other benchmark datasets (OntoNotes and PreCo) and those trained on our newly created coreference annotations in LitBank.

### 6.1. Gold mentions

OntoNotes, PreCo and LitBank all differ in their annotations in two important ways: not only in the individual coreference decisions (where, for example, OntoNotes and LitBank do *not* link two generic common nouns, while PreCo does), but also in the core annotation of mentions—while OntoNotes does not annotate singleton mentions (those entities that are not coreferent with any other), PreCo and LitBank both do. Additionally, while OntoNotes and PreCo both capture unrestricted entity coreference (annotating all entities regardless of their entity type), LitBank only annotates those that fall in the ACE entity categories PER, FAC, LOC, GPE, ORG, and VEH. In order to assess the potential for models trained on these datasets to accurately capture this kind of coreference, we first evaluate on gold mention spans in LitBank.

For each dataset, we train a neural model based on Lee et al. (2017). While Lee et al. (2017) is an end-to-end mention-ranking model that jointly performs mention identification and linking, we constrain it to fixed mention spans and use it solely for mention ranking, greedily selecting the most likely antecedent for each mention in a document (or a null antecedent that begins a new coreference chain with that mention). We also differ by representing each token in a sentence not by static Glove/Turian word embeddings and character convolution, but through its BERT token representation (Devlin et al., 2019). While BERT uses WordPiece tokenization (Wu et al., 2016), we average together the individual WordPiece tokens for a given word to form its token representation, as in Sims et al. (2019).

This model is a bidirectional LSTM that generates output  $x_i$  for each token at position  $i$ . A mention  $m$  spanning sentence positions  $[start, end]$  is represented as the concatenation of  $x_{start}, x_{end}$ , the output of a learned attention mechanism over  $[x_{start}, \dots, x_{end}]$ , and features expressing the width of the mention span and whether the span falls within a quotation (each feature is embedded in its own representation space). Given two mention representations  $g_i$  and  $g_j$ , the score of their linking is a feedforward network over the concatenation of  $g_i, g_j$ , the elementwise product of  $g_i$  and  $g_j$ , and a feature function scoped over the two mentions; we use features expressing the distance between mentions (in mentions and sentences) and whether one mention is nested within the other. This model is trained to maximize the marginal log-likelihood of all antecedents in the correct coreference chain for each mention; during prediction, we proceed from the beginning of the document to the end, and greedily select the single highest-scoring antecedent (or the null antecedent) for each mention, and define a coreference chain as the transitive closure of all such links.

To summarize, this model is identical with Lee et al. (2017) with the following exceptions: 1.) We use BERT contextual

representations instead of static word vectors and a subword character CNN (to make use of advancements in representation learning); 2.) we train and predict conditioning on mention boundaries (in order to separate the core task of mention linking from mention identification, given the differences in mention boundaries in the three datasets; 3.) We omit author and genre information from training (as these are relevant only in OntoNotes and not in PreCo or LitBank); 4.) We only consider antecedents within 20 mentions for pronouns and 300 mentions for proper noun phrases and common noun phrases (given the observations on distance distributions above).

Aside from these differences, we preserve the same core hyperparameter choices in Lee et al. (2017): LSTM size of 200; the feedforward network is comprised of two 150-dimensional layers; each feature is embedded in a learned 20-dimensional space; and we use ADAM (Kingma and Ba, 2014) for learning, decaying the learning rate by 0.1% each 100 steps. Each training source has a training and development split; we train each model until there is no improvement for 10 epochs on its development set, saving the best-performing model on that development data. We test performance on literary data by training on three different datasets: OntoNotes, PreCo and our literary data; we evaluate on the entirety of the literary annotations (using cross-validation where needed, as described below), excluding copula and apposition links. Table 3 presents the results of this evaluation using F-scores from  $B^3$  (Bagga and Baldwin, 1998), MUC (Vilain et al., 1995) and CEAF $_{\phi_4}$  (Luo, 2005), along with their average.

**OntoNotes.** The training set of OntoNotes contains a total of 1.3M tokens; when evaluated on its own test set, we see performance comparable to that reported in Lee et al. (2017) for gold mentions, along with an expected drop in performance for not using speaker or genre metadata (83.2 vs. 85.2). When evaluated on LitBank, however, this performance drops 10.3 points to 72.9.

**PreCo.** The training set of PreCo contains an order of magnitude more data, with a total of 12.2M tokens. This increase in training size relative to OntoNotes translates into an improvement in overall accuracy: when evaluated on LitBank, we see an average F1 score of 78.8.

**LitBank.** To evaluate on LitBank, we perform a 10-fold cross-validation—training a model on 80% of the data, assessing early stopping using a development set of 10% of the data, and then evaluating that trained model on the remaining held-out 10%—over all ten partitions. Such a model achieves an average F-score of 79.3, indistinguishable from a model trained on PreCo but substantially better than one trained on OntoNotes.

Training source	$B^3$	MUC	CEAF $_{\phi_4}$	Average
OntoNotes	66.9	85.7	65.9	72.9
PreCo	73.8	88.4	74.3	78.8
LitBank	72.6	88.5	76.7	79.3

Table 3: Coreference resolution performance on gold mentions.

A natural reason why OntoNotes may underperform on this

dataset is the mismatch between annotation guidelines—especially the difference between the choice by PreCo and LitBank to annotate singleton mentions and OntoNotes not to. To assess this possibility, we can leverage the  $B^3$  coreference metric to interrogate performance on a subset of the data. Precision (P) and recall (R) in  $B^3$  is calculated by finding the overlap between a gold coreference chain  $\mathcal{G}_i$  and system coreference chain  $\mathcal{S}_i$  associated with each mention  $i$  over all  $N$  mentions:

$$P = \frac{1}{N} \sum_{i=1}^N \frac{|\mathcal{G}_i \cap \mathcal{S}_i|}{|\mathcal{S}_i|} \quad R = \frac{1}{N} \sum_{i=1}^N \frac{|\mathcal{G}_i \cap \mathcal{S}_i|}{|\mathcal{G}_i|}$$

We can adapt this to examine non-singleton performance by ranging only over mentions that are not singletons in the gold data; if we exclude true singleton mentions from our evaluation, the performance for all three training sources drops (since singletons are comparatively easier to identify)—with gold mention boundaries, PreCo  $B^3$  drops from 73.8 to 69.7, OntoNotes  $B^3$  drops from 66.9 to 63.3, and LitBank  $B^3$  drops from 72.6 to 67.7. OntoNotes drops less than the other methods—attesting to some impact of the singleton annotation difference—but still lags behind the other two sources of data on this metric.

## 6.2. Predicted mentions

To test real-world performance when not given access to gold mention boundaries, we build a mention identifier trained on the mention boundaries in our annotated literary data, and assess coreference resolution performance when using these predicted mentions at test time.

### 6.2.1. Mention identification

We train a mention identification system with a layered BiLSTM-CRF of Ju et al. (2018) used in Bamman et al. (2019); we adapt it here by replacing the static word embeddings in the original work with BERT contextual embeddings. In a tenfold cross-validation on our annotated data to predict the mention spans, we find an F-score for mention detection to be 89.1. Adding a prediction for whether an entity is a proper noun phrase, common noun phrase or pronoun only decreases the F-score to 88.3; adding a prediction for the entity class (person, organization, location, facility, geo-political entity, or vehicle) yields a performance of 87.3.

Task	Precision	Recall	F
Mention span detection	90.7	87.6	89.1
+ PROP/NOM/PRON	90.2	86.5	88.3
+ Entity class	89.2	85.5	87.3

Table 4: Mention identification performance.

### 6.2.2. Coreference resolution performance

To assess performance on predicted mentions, we carry out another ten-fold cross-validation: for one partition of the data into 80% training data, 10% development data, and 10% test data, we use a mention identification model trained on the training split to predict mention boundaries in the test split, and then train our neural coreference system on that same train split to make coreference decisions on the predicted mentions in the test split (again using the

development set to assess early stopping). Table 5 presents the results of this evaluation. While performance naturally degrades as a function of the predicted mentions, we see the same overall rank between training sources as found in table 3: LitBank (68.1F) and PreCo (67.6F) are indistinguishable but both perform substantially better than OntoNotes.

Training source	$B^3$	MUC	CEAF $_{\phi_4}$	Average
OntoNotes	57.7	81.2	49.7	62.9
PreCo	63.5	84.2	55.1	67.6
LitBank	62.7	84.3	57.3	68.1

Table 5: Coreference resolution performance on predicted mentions.

## 7. Conclusion

We present in this work a new dataset of coreference annotations for 210,532 tokens drawn from 100 different literary texts, to allow the analysis of coreference performance in a wide range of literary styles and authors, and to train coreference resolution systems on literary data in English. As more and more work in literary analysis makes use of computational methods to explore the persistence of entities in text—from characters (Piper, 2018; Underwood, 2019) to objects (Tenen, 2018)—having reliable estimates of the quality of different coreference methods is critical for assessing their impact on downstream measurements.

In addition to the core contribution of this dataset and evaluation, there are three takeaways worth highlighting: the first is the relatively high performance of PreCo in cross-domain coreference resolution. While PreCo comes from a separate domain of school examinations, it is able to match the performance of a smaller dataset of in-domain coreference annotations; while the literary annotations are able to achieve comparable performance with two orders of magnitude less data, PreCo illustrates the power of large annotated datasets to work well across a potential variety of domains. Second, we illustrate the behavior of entities across long spans of text—demonstrating that entities tend to be bimodal (either active over very short timespans or very long ones), and frequent entities tend to exhibit bursty behavior over the course of long documents. Third, we illustrate the differing antecedent behavior of different mention types, showing that pronouns in particular can be linked to an antecedent within nine mentions 95% of the time. We hope this work can spur future work on literary coreference, since the characteristics of literature—the relatively long length of documents, the centrality of relatively few major entities, the difficult metaphysical and epistemological questions of identity and the revelation of knowledge—can potentially provide a unique vantage point on the problem of coreference in general. This data is freely available at <https://github.com/dbamman/litbank>; code to support this work can be found at <https://github.com/dbamman/lrec2020-coref>.

## 8. Acknowledgments

The research reported in this article was supported by an Amazon Research Award and by resources provided by NVIDIA and Berkeley Research Computing.



## 9. Bibliographical references

- Agarwal, A., Corvalan, A., Jensen, J., and Rambow, O. (2012). Social network analysis of alice in wonderland. In Proceedings of the NAACL-HLT 2012 Workshop on Computational Linguistics for Literature, pages 88–96, Montréal, Canada, June. Association for Computational Linguistics.
- Bagga, A. and Baldwin, B. (1998). Algorithms for scoring coreference chains. In The First International Conference on Language Resources and Evaluation Workshop on Linguistics Coreference, volume 1, pages 563–566. Granada.
- Bamman, D., Underwood, T., and Smith, N. A. (2014). A Bayesian mixed effects model of literary character. In Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 370–379, Baltimore, Maryland, June. Association for Computational Linguistics.
- Bamman, D., Popat, S., and Shen, S. (2019). An annotated dataset of literary entities. 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 2138–2144, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Chen, H., Fan, Z., Lu, H., Yuille, A., and Rong, S. (2018). PreCo: A large-scale dataset in preschool vocabulary for coreference resolution. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pages 172–181, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Clark, K. and Manning, C. D. (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, August. Association for Computational Linguistics.
- Cohen, K. B., Lanfranchi, A., Choi, M. J.-y., Bada, M., Baumgartner, W. A., Panteleyeva, N., Verspoor, K., Palmer, M., and Hunter, L. E. (2017). Coreference annotation and resolution in the Colorado Richly Annotated Full Text (CRAFT) corpus of biomedical journal articles. *BMC Bioinformatics*, 18(1):372, Aug.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (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, June. Association for Computational Linguistics.
- D’Souza, J. and Ng, V. (2012). Anaphora resolution in biomedical literature: A hybrid approach. In Proceedings of the ACM Conference on Bioinformatics, Computational Biology and Biomedicine, BCB ’12, pages 113–122, New York, NY, USA. ACM.
- Elson, D. K., Dames, N., and McKeown, K. R. (2010). Extracting social networks from literary fiction. In Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, ACL ’10, pages 138–147, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Gallois, A. (2016). *The Metaphysics of Identity*. Routledge.
- Gasparin, C., Karamanis, N., and Seal, R. (2007). Annotation of anaphoric relations in biomedical full-text articles using a domain-relevant scheme. In Proceedings of DAARC, volume 2007. Citeseer.
- Ghaddar, A. and Langlais, P. (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, May. European Language Resources Association (ELRA).
- Guha, A., Iyyer, M., Bouman, D., and Boyd-Graber, J. (2015). Removing the training wheels: A coreference dataset that entertains humans and challenges computers. In Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 1108–1118, Denver, Colorado, May–June. Association for Computational Linguistics.
- Hovy, E., Marcus, M., Palmer, M., Ramshaw, L., and Weischedel, R. (2006). OntoNotes: the 90% solution. In Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers, NAACL-Short ’06, pages 57–60, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Jannidis, F., Reger, I., Krug, M., Weimer, L., Macharowsky, L., and Puppe, F. (2016). Comparison of methods for the identification of main characters in German novels. In *Digital Humanities*.
- Ju, M., Miwa, M., and Ananiadou, S. (2018). A neural layered model for nested named entity recognition. 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 1446–1459, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Kantor, B. and Globerson, A. (2019). Coreference resolution with entity equalization. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 673–677, Florence, Italy, July. Association for Computational Linguistics.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. <http://arxiv.org/abs/1412.6980>.
- Kraicer, E. and Piper, A. (2018). Social characters: The hierarchy of gender in contemporary English-language fiction. *Cultural Analytics*.
- Krug, M., Puppe, F., Jannidis, F., Macharowsky, L., Reger, I., and Weimar, L. (2015). Rule-based coreference resolution in German historic novels. In Proceedings of the Fourth Workshop on Computational Linguistics for Literature, pages 98–104, Denver, Colorado, USA, June. Association for Computational Linguistics.
- Krug, M., Weimer, L., Reger, I., Macharowsky, L., Feldhaus, S., Puppe, F., and Jannidis, F. (2017). Description of a corpus of character references in German nov-

- els - DROC [Deutsches Roman Corpus]. DARIAH-DE Working Papers No. 27.
- Lee, J. and Yeung, C. Y. (2012). Extracting networks of people and places from literary texts. In Proceedings of the 26th Pacific Asia Conference on Language, Information, and Computation, pages 209–218, Bali, Indonesia, November. Faculty of Computer Science, Universitas Indonesia.
- Lee, K., He, L., Lewis, M., and Zettlemoyer, L. (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, September. Association for Computational Linguistics.
- Lee, K., He, L., and Zettlemoyer, L. (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, June. Association for Computational Linguistics.
- Lestari, V. A. (2018). Building effective representations for domain adaptation in coreference resolution. Ph.D. thesis, UT-Austin.
- Luo, X. (2005). On coreference resolution performance metrics. In Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing, HLT '05, pages 25–32, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Moosavi, N. S. and Strube, M. (2017). Lexical features in coreference resolution: To be used with caution. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pages 14–19, Vancouver, Canada, July. Association for Computational Linguistics.
- Moosavi, N. S. and Strube, M. (2018). Using linguistic features to improve the generalization capability of neural coreference resolvers. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pages 193–203, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Nguyen, N., Kim, J.-D., and Tsujii, J. (2011). Overview of BioNLP 2011 protein coreference shared task. In Proceedings of BioNLP Shared Task 2011 Workshop, pages 74–82, Portland, Oregon, USA, June. Association for Computational Linguistics.
- Nunberg, G. (1984). Individuation in context. In Proceedings of the Second West Coast Conference on Formal Linguistics.
- Piper, A., Algee-Hewitt, M., Sinha, K., Ruths, D., and Vala, H. (2017). Studying literary characters and character networks. In DH 2017.
- Piper, A. (2018). Enumerations. University of Chicago Press.
- Pradhan, S., Ramshaw, L., Marcus, M., Palmer, M., Weischedel, R., and Xue, N. (2011). CoNLL-2011 shared task: Modeling unrestricted coreference in OntoNotes. In Proceedings of the Fifteenth Conference on Computational Natural Language Learning: Shared Task, pages 1–27, Portland, Oregon, USA, June. Association for Computational Linguistics.
- Pradhan, S., Moschitti, A., Xue, N., Uryupina, O., and Zhang, Y. (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, July. Association for Computational Linguistics.
- Recasens, M., Hovy, E., and Martí, M. A. (2010). A typology of near-identity relations for coreference (NIDENT). In Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10), Valletta, Malta, May. European Language Resources Association (ELRA).
- Recasens, M., Hovy, E., and Martí, M. A. (2011). Identity, non-identity, and near-identity: Addressing the complexity of coreference. *Lingua*, 121(6):1138 – 1152.
- Recasens, M., de Marneffe, M.-C., and Potts, C. (2013). The life and death of discourse entities: Identifying singleton mentions. In Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 627–633. Association for Computational Linguistics.
- Reiter, N. and Frank, A. (2010). Identifying generic noun phrases. In Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, ACL '10, pages 40–49, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Roesiger, I., Schulz, S., and Reiter, N. (2018). Towards coreference for literary text: Analyzing domain-specific phenomena. In Proceedings of the Second Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature, pages 129–138, Santa Fe, New Mexico, August. Association for Computational Linguistics.
- Schäfer, U., Spurk, C., and Steffen, J. (2012). A fully coreference-annotated corpus of scholarly papers from the ACL anthology. In Proceedings of COLING 2012: Posters, pages 1059–1070, Mumbai, India, December. The COLING 2012 Organizing Committee.
- Sims, M., Park, J.-H., and Bamman, D. (2019). Literary event detection. In ACL.
- Stenertorp, P., Pyysalo, S., Topić, G., Ohta, T., Ananiadou, S., and Tsujii, J. (2012). BRAT: A web-based tool for NLP-assisted text annotation. In Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics, EACL '12, pages 102–107, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Su, J., Yang, X., Hong, H., Tateisi, Y., and Tsujii, J. (2008). Coreference resolution in biomedical texts: a machine learning approach. In Dagstuhl Seminar Proceedings. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- Subramanian, S. and Roth, D. (2019). Improving generalization in coreference resolution via adversarial training. In Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (\*SEM 2019), pages

- 192–197, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Sudhahar, S. and Cristianini, N. (2013). Automated analysis of narrative content for digital humanities. In *International Journal of Advanced Computer Science*.
- Tenen, D. Y. (2018). Toward a computational archaeology of fictional space. *New Literary History*.
- Underwood, T., Bamman, D., and Lee, S. (2018). The transformation of gender in English-language fiction. *Cultural Analytics*.
- Underwood, T. (2019). *Distant Horizons: Digital Evidence and Literary Change*. University of Chicago Press.
- Vala, H., Jurgens, D., Piper, A., and Ruths, D. (2015). Mr. Bennet, his coachman, and the Archbishop walk into a bar but only one of them gets recognized: On the difficulty of detecting characters in literary texts. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 769–774, Lisbon, Portugal, September. Association for Computational Linguistics.
- Vala, H., Dimitrov, S., Jurgens, D., Piper, A., and Ruths, D. (2016). Annotating characters in literary corpora: A scheme, the CHARLES tool, and an annotated novel. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 184–189, Portorož, Slovenia, May. European Language Resources Association (ELRA).
- Vilain, M., Burger, J., Aberdeen, J., Connolly, D., and Hirschman, L. (1995). A model-theoretic coreference scoring scheme. In *Proceedings of the 6th Conference on Message Understanding, MUC6 '95*, pages 45–52, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Walker, C., Strassel, S., Medero, J., and Maeda, K. (2006). ACE 2005 multilingual training corpus. LDC.
- Wiseman, S., Rush, A. M., and Shieber, S. M. (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, San Diego, California, June. Association for Computational Linguistics.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, L., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., and Dean, J. (2016). Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144.
- Yang, X., Zhou, G., Su, J., and Tan, C. L. (2005). Improving noun phrase coreference resolution by matching strings. In *Proceedings of the First International Joint Conference on Natural Language Processing, IJCNLP'04*, pages 22–31, Berlin, Heidelberg. Springer-Verlag.
- Zeldes, A. (2017). The GUM corpus: Creating multi-layer resources in the classroom. *Language Resources and Evaluation*, 51(3):581–612.