

Opinion Piece: Can we Fix the Scope for Coreference? Problems and Solutions for Benchmarks beyond OntoNotes

Amir Zeldes

*Georgetown University
Department of Linguistics*

AMIR.ZELDES@GEORGETOWN.EDU

Editor: Massimo Poesio

Submitted 04/2021; Accepted 01/2022; Published online 04/2022

Abstract

Current work on automatic coreference resolution has focused on the OntoNotes benchmark dataset, due to both its size and consistency. However many aspects of the OntoNotes annotation scheme are not well understood by NLP practitioners, including the treatment of generic NPs, noun modifiers, indefinite anaphora, predication and more. These often lead to counterintuitive claims, results and system behaviors. This opinion piece aims to highlight some of the problems with the OntoNotes rendition of coreference, and to propose a way forward relying on three principles: 1. a focus on semantics, not morphosyntax; 2. cross-linguistic generalizability; and 3. a separation of identity and scope, which can resolve old problems involving temporal and modal domain consistency.

Keywords: coreference, annotation, guidelines, anaphora, referring expressions, predication, scope, multilingual, OntoNotes

1. Introduction

Coreference resolution is the task of delineating and grouping together referring expressions in text so that spans of text referring to the same discourse entity are clustered together. The past decade has seen remarkable improvements in the consistency and performance of coreference resolution, first through the creation of large (>1 M tokens) benchmark data in OntoNotes (Hovy et al. 2006; Weischedel et al. 2012; hence ON), and then through use of that dataset in the CoNLL shared task on coreference resolution (Pradhan et al., 2011) and the development of evaluation metrics (Pradhan et al. 2014; see Moosavi and Strube 2016 for criticism). With the advent of end-to-end neural approaches to coreference resolution (Lee et al., 2017), we have seen scores on coreference resolution rise from the mid-50s at the 2011 shared task to current SOTA scores around 80 points when gold speaker information is provided (Joshi et al., 2019; Wu et al., 2020), as evaluated by the standard metrics on the ON test set.

At the same time, coreference resolution using the ON scheme as a target has been plagued by a number of issues: the lack of annotation of singletons (entities mentioned only once) has led to systems conflating referentiality recognition (whether an expression in fact refers to some entity)

I would like to thank Anna Nedoluzhko, Massimo Poesio, Sameer Pradhan, Nathan Schneider and the anonymous reviewers for valuable comments and discussions about previous versions of this paper; the usual disclaimers apply.

with coreference recognition (whether a referring expression is mentioned more than once, see Lee et al. 2017; Zhang et al. 2018); omission of annotation of indefinite anaphors has led to counter-intuitive system behaviors (e.g. repeated mention of specific indefinites and generics is ignored; see Section 2.1 and the Appendix for examples); and idiosyncratic handling of some constructions (especially copula predication, compounding) has meant that even when systems perform well in predicting ON-like annotations, the results are often less informative than they could be for practical applications (see Section 2.4). Additionally, because the ON guidelines for English depend on morphosyntactic constructs such as indefinite NPs, copula predication and compounding, they are not readily applicable to other languages, leading to rather different practices in corpora for languages other than English, including in the Arabic and Chinese sections of the OntoNotes corpus itself.

I would like to start by saying that I have been a passionate admirer and intensive user of OntoNotes for a long time, but I also suspect that dissatisfaction with some aspects of its annotation scheme is widespread: this is evidenced by the emergence of the Coreference Resolution Beyond OntoNotes workshops since 2016 (CORBON, and later CRAC, ‘Computational Models of Reference, Anaphora and Coreference’), the continuing creation of resources with different annotation schemes (cf. Stoyanov et al. 2009; Zeldes and Zhang 2016; see Poesio et al. 2016 for an overview), the focus of the 2021 coreference shared task on non-ON data, and most recently the launching of a Universal Anaphora (UA) initiative, aimed at bringing some consistency into the growing variety of divergences between resources.¹ However, I suspect that while we as a research community can arrive at a consensus about this dissatisfaction, it will be much harder to reach a consensus about what a solution might look like, and how it could be attained feasibly. My goal in writing this overview is therefore first to outline some of the reasons why we need to do something, and my audience is, in the first instance, the community working on language resources for coreference, and in the second instance, NLP researchers, and especially grad students, who are entering the field and may get the impression that the ON standard represents a final end point of what we want from coreference resolution. The suggestions below probably form a continuum between more and less controversial, and their ranking on different researchers’ wish lists will differ as well based on their respective goals, but it’s a discussion we can only have if the commonalities and differences between coreference datasets and task definitions are laid out clearly.²

2. The 90% and 10% in OntoNotes

The development and release of ON were landmark achievements for a number of annotation types covered by the corpus, but especially transformative for coreference resolution. One of the declared goals of the project, as presented in the Hovy et al. paper and its title, was to reach a 90% solution: agreement scores for annotations had to be in the 90s, casting consistency as a necessary condition for any guidelines developed for the corpus. For coreference, this meant focusing on some of the easier phenomena to agree on, such as antecedents for pronominal anaphora, while more tricky cases, such as indefinite anaphors, were left unannotated – not because they were deemed disinteresting or unuseful, but because they were a lower priority target (Sameer Pradhan, p.c.). This decision led to high agreement, but also to a number of unexpected results.

¹See <http://universalanaphora.org/> and Nedoluzhko et al. (2022).

²For readers who prefer to start with worked examples, see the Appendix for some analyses contrasting the ON view of coreference with a more unrestricted annotation scheme.

In this section I will survey some of the most surprising cases for lay users or researchers new to coreference resolution; for quite a few cases, I am certain that the majority of people using off-the-shelf coreference tools are not well aware of the specifics or consequences, and many people working with ON data directly are also not likely to know about all. In the following section I will argue that for all of these, semantic reference to the same entity, rather than morphosyntactic environments, should guide the coreference decision.

2.1 Indefinites and generics

ON guidelines prohibit “generic, underspecified, and abstract nominal mentions” (BBN Technologies, 2007, 5) from being linked to other such mentions, meaning that they can only be referred back to by pronouns or definite NPs.³ Although indefinite NPs can indicate a variety of semantic or pragmatic functions (Bhatia et al., 2014), all “indefinite noun phrases, which begin with the indefinite article (a, an)” (BBN Technologies, 2007, 6) as well as bare NPs (bare plurals and singulars with no article) are considered generic by definition in ON. Thus the following are *not* annotated in the corpus in any way:

- (1) [Program trading]_i is “a racket,” ... [program trading]_i creates deviant swings
(ON, wsj_0121, no coreference in the corpus)

This indirectly implies that coreference resolution systems pairing a pronoun or definite mention with the first NP in such a chain will be penalized, and that systems must learn in an environment in which ‘it’ + ‘program trading’ is once a positive match and once a negative one, despite the fact that it is semantically the same ‘it’, which in fact refers to the generic concept of program trading. The same guideline applies to generic pronouns, as in the following example, which is also left unannotated in ON.

- (2) [you]_i couldn’t start unless [you]_i knew that the replacement heart would make it to the operating room (BBN Technologies, 2007, 5)

In this case, systems must again learn that the usual coreference between two instances of ‘you’ in the same sentence does not apply, even though semantically, I would again argue that both pronouns do in fact refer to the same (albeit generic) thing – anyone who could have started a heart replacement in this context. From a set-theoretical perspective it may even be possible to argue that that set (people who could have started heart replacements at the time) is completely specified, and non-generic (or at least, we could continue the text in a way which explicitly states the names of the relevant people).

The problems with the abstract/generic guideline as a foundation for universal, cross-linguistic coreference annotation are numerous:

³An anonymous reviewer has asked about the sense of ‘referring’ or ‘referentiality’ intended here. The term is unfortunately murky: on the one hand, theoretical literature going back to at least Reinhart (1982) has referred to cognitive models such as ‘file cards’ per entity (see Krifka 2008 for discussion), which can be evoked by language, but which are empirically hard to define or observe. On the other hand the practical needs of coreference resolution in NLP have often equated referentiality with potential for back-reference. In many partial implementations of ‘tricky’ coreference, this has amounted to including usually ignored expressions if they are referred back to, such as verbal markables for discourse deixis (‘she visited... the visit’), or limited markup of coordinations when they are mentioned again with a plural pronoun (‘Kim and Yun ... they’).

1. Indefinite mentions are often neither generic nor underspecified or abstract (“[*participants*] comprised [15 women and 10 men]”)
2. It is not immediately obvious that we should not want coreference information even for those mentions which are generic, abstract, etc. (e.g. “*program trading*”), especially given that we do resolve pronouns referring back to them
3. In context it is often hard to know whether pronominal mentions are generic, and even when they are, they can still form multiple distinct clusters
4. Many languages do not have widespread articles which could be used to identify generic mentions, even if we agreed that all indefinites should be considered generic

To illustrate this, we can consider the following examples from the GUM corpus (Zeldes, 2017), a freely available English coreference data set with texts from 12 written and spoken text types, which does not exclude indefinite anaphors:⁴

- (3) Marbles is the first social media star to have [a wax figure]_i displayed in Madame Tussauds ... In 2015 , Marbles unveiled [a wax figure of herself]_i at Madame Tussauds (GUM_bio_marbles)
- (4) I have [a mini-MMPI]_i ... I have [a chart that I’ll go through]_i (GUM_interview_dungeon)

In (3), the wax figure is a specific, unique physical object, mentioned twice as an indefinite NP in the same document. In (4), the ‘chart’ is referring in context to the exact same entity as the mini-MMPI (a ‘Minnesota Multiphasic Personality Inventory’) – without the coreference information, a tool or researcher accessing the data automatically would not know that the MMPI was a chart, and vice versa, despite such information sharing being an obvious application of coreference resolution.⁵

Generic pronouns too can often serve different purposes and belong to different, meaningful clusters:

- (5) [You]₁ feel like [you]₁’re prepared , [you]₁’re in a , [you]₂ know , in a relationship ... [you]₂ know (GUM_vlog_pregnant)

In (5), which discusses feeling prepared for a pregnancy, all instances of ‘you’ are arguably generic, but they can be grouped into two sets: the first set refers to some person who might be considering a pregnancy (thereby implicitly ruling out those who cannot become pregnant), while the second set is used exclusively as a pragmatic discourse marker (Östman, 1981) to refer to a generic audience and solicit presupposed agreement. It is therefore clear that in context, the person who may be prepared is the same person filling the subject role for ‘feel’, while ‘you’ in ‘you know’ can refer to anyone who may be listening, which forms a different set of individuals (the original utterance comes from a vlog, and the audience is therefore underspecified even for the speaker).

Finally, the issue of cross-linguistic applicability surfaces already in the ON guidelines for Chinese, which cannot refer to indefinite articles as a criterion for identifying genericity, since Chinese

⁴See the corpus website (<https://corpling.uis.georgetown.edu/gum>) for example sources, and the Appendix for longer examples.

⁵An anonymous reviewer has cautioned that this should not imply that disjoint indefinite mentions should corefer, as in ‘I always owned [a dog]_{generic}. For my 6th birthday, I got [a dog]_x. When it died we got [a dog]_y from the animal sanctuary.’ I fully agree that there are three distinct ‘dog’ entities in this case, where x and y are subsets of the generic ‘dog’, and therefore not coreferring.

does not use articles as in English.⁶ The guidelines instead refer to specific examples: “死刑 (capital punishment), 世界 (world – as distinct from “the world” meaning the planet Earth), 社会 (society) are considered generic nouns” (BBN Technologies, 2008, 8). The ON Chinese corpus itself, however, is not consistent about what cases are excluded as generic, and even these specified nouns are regularly annotated as coreferring in contexts which would not be considered definite or specific in English. For ‘society’, these include 44 coreferring instances, as in (6).

- (6) 我们能不能发展的快一些、好一些，实现经济快速发展和 [社会]_i 全面进步，并且保持 [社会]_i 稳定，十分重要 (ON, cnr_0016)
Wǒmen néng bùnéng fāzhǎn de kuài yīxiē, hǎo yīxiē, shíxiàn jīngjì kuàisù fāzhǎn hé [shèhuì] quánmiàn jìnbù, bìngqiě bǎochí [shèhuì] wěndìng, shí fèn zhòngyào
 Whether our development can progress faster and better to make it possible for the economy to grow quickly and for [society]_i to make progress across all metrics as well as to maintain the stability of [society]_i is very crucial

Rather than concluding that these cases are annotation errors, I would like to argue below in line with previous work showing the salience and coreferring potential of a range of indefinite phrases (see e.g. Kunz 2009, 295–300) that it is reasonable to annotate them, just as the recurring annotations suggest that annotators regarded them as coreferring.⁷

2.2 Compound modifiers

Compound modifiers (or noun-modifying nouns) in English, as in many languages with similar ‘bare’ adnominal modification, often have low referential content, and are known to resist pronominalization (i.e. they are ‘anaphoric islands’, Postal 1969), making them attractive to exclude categorically. For example, as a compound nominal premodifier, it is difficult to imagine pronominal reference to ‘animals’ introduced by a compound ‘animal hunters’ in (7). Because of this often limited referential potential, and the desire to avoid disagreements, a conservative annotation scheme may prioritize simplicity and consistency by prohibiting compound modifiers from entering into coreference relations. At the same time, repeated mention of proper nouns in the same position seems normal, and not indefinite or non-referential, as in (8):

- (7) *[Animal]_i hunters tend to like [them]_i. (Postal, 1969, 230)
 (8) The [Hong Kong]_i government’s jurisdiction is the [Hong Kong]_i Special Administrative Region

ON makes an exception for proper noun modifiers (BBN Technologies, 2007, 3), but excludes common noun cases like (9), which are left unannotated despite subsequent definite reference.

⁶The same is true of many languages; for example, Slavic language coreference datasets have included generic NPs, which cannot be distinguished based on articles (see e.g. Nedoluzhko et al. 2009 for Czech).

⁷One reviewer pointed out that we may still want to use articles in guidelines for languages that do distinguish them, as this can raise inter-annotator agreement, which may be the more important factor for some applications. I certainly do not mean that any reference to articles in guidelines for such languages should be forbidden; but conversely, ruling out ‘all indefinite anaphors’ in coreference datasets a priori, based on form alone, will not produce comparable data across languages, and as I hope this section shows, this would make us lose important information for English as well.

- (9) small investors seem to be adapting to greater [stock market]_i volatility ... Glenn Britta ... is “factoring” [the market’s]_i volatility “into investment decisions.” (ON, wsj_0121)

However it is by now well known that compound modifiers are not strict islands even to pronominal anaphora (Ward et al., 1993), and corpus data provides plentiful counter examples, leading to many of the same issues as those involved with generics (or more generally, indefinites): noun modifiers can be specific and can be referred back to with definite expressions as in (9), and even when they are truly generic, we may want to know about their coreference. In (10), the indication that cinnamon is a spice (and not a brand, a person’s name, or something else) is made explicit via coreference, and it is not clear to me why we should exclude such cases from the target gold standard for annotation and NLP tools.

- (10) [Cinnamon]_i basil really does smell like [the sweet spice]_i (GUM_what_basil)

As before, the restriction on compound modifiers in English is problematic from a cross-linguistic perspective, and in fact, construct-state modifiers in the ON Arabic coreference annotations, which are very similar to English compound modifiers (e.g. modifiers resist bearing articles, the translation of (7) would be ungrammatical), are not restricted in the same way, as seen in (11) with a semantically specific modifier:

- (11) محاكمة ضباط روس بتهمة الاهمال ... محاكمة ... لثلاثة ضباط روس (ON, ann_0006)
muḥākamatu ḍubbātin rūsi bituhmati l-ihmāli... muḥākamatun... li-talātatin ḍubbātin rūsi
 [Russian Officer]_i Trial on Charges of Negligence... a trial... for [three Russian officers]_i

Here the Russian officers in the ‘Russian officer trial’ are exactly the same set of ‘three officers’. Such examples with specific indefinite NPs create the most obvious motivation for including modifier nouns, but the same construction and coreference appear equally with non-specific/generic referents. For parallel corpora and applications of coreference resolution in machine translation too, we can note that what is expressed by such modifiers in one language may be expressed by a full NP in another, meaning that excluding them will lead to more cross-linguistic divergence (see Novák and Nedoluzhko 2015, 31 for example for English and Czech).

The inconsistency between the annotations of modifier nouns in the three ON languages is thus not only odd from a theoretical perspective; it can cause problems from a practical one as well if we try to develop multilingual applications relying on automatic coreference resolution, only to find out that even systems trained on ON itself behave fundamentally differently for different languages.

2.3 Nesting

Nested coreference, i.e. a mention with a coreferring mention within its own span, may seem odd, but is in fact rather frequent, and attested in this very sentence. ON does allow nested coreference, but excludes cases without subsequent mention, leading to another difficult learning task for systems (the non-markable status of such cases hinges on later sentences), and unresolved pronouns in text, as in (12), which remains unannotated in OntoNotes, presumably because outside of the larger phrase’s boundaries, it appears to be a singleton.⁸

⁸In fact, some systems have relied on this constraint, sometimes called ‘i-within-i’ due to the co-index *i*, to rule out certain kinds of match candidates, see Lee et al. (2013, 896).

(12) [an elusive sheep with a star on [its]_i back]_i (ON, wsj_0037)

A further complication in ON concerns nested dates, in which only the top-level expression is considered referential. For example, if a year is mentioned by itself, it may corefer with other mentions, but if it is part of a date, then it may not (again, the example is unannotated in ON):

(13) what are the opportunities for new developments in the wake of the [1999]_i handover? On December 19, [1999]_i, the eve of Macau’s transfer of sovereignty... (ON, ectb_1001)

In (13), we are unable to discover that the year 1999 in the first sentence is the same year 1999 as in the second, since annotation is ruled out by the ON *form-based* guideline on nested dates. Some readers may think this is not so bad, since we can surely analyze both dates and discover that the year is the same; but I would argue we need to keep in mind that date detection is not trivial (1999 could be a price, or something else), and that these two cases should co-refer based on a *semantic* criterion (they refer to the same thing in the world).

Moreover, for system development it is problematic to create examples where algorithms will learn that date expressions are atomic, since we can have overt underspecified anaphoric expressions targeting subspans of dates:

(14) between March 18, [2005]_i and May 7 of [the same year]_i (ON, a2e_0016)

Example (14) actually *is* annotated in ON, in contradiction with the guidelines, perhaps because of how counterintuitive their effect would be in this case and others like it (this example is by no means unique in ON, and others are left unannotated as intended). Without coreference in (14), we have no chance of discovering that “2005” is “the same year” in which “May 7” takes place. Such information can be crucial for identifying DCT (Document Creation Time, Ray et al. 2018) and for Event Timeline extraction (e.g. for clinical events, Nikfarjam et al. 2013), among other applications. However the main point from my perspective is that this behavior is unexpected from a semantic point of view, and is narrowly tied to specific forms, not meanings.

2.4 Predication

Judging by discussions with some of my colleagues, perhaps the most controversial opinion I can express in this paper is that the omission of predication from mainstream coreference resolution datasets and systems is a mistake.⁹ Following an early period in the development of coreference datasets (notably ACE, Doddington et al. 2004) in which most nominal predicates were considered to corefer to their subjects (i.e. “Kim is a teacher” \Rightarrow “Kim” $\xleftrightarrow{\text{coref}}$ “a teacher”), criticism expressed for example by van Deemter and Kibble (2000) and Zaenen (2006) led to ON rejecting coreference in predication altogether. In this section I would like to suggest not only that ON’s guidelines are an overreaction to the original criticism, but also that omission is neither a sufficient nor a necessary remedy for the problems that led to its rejection, an outcome left open by van Deemter and Kibble themselves.

The core of the problem with predication has been expressed as “change over time” (Hirschman and Chinchor 1998, 11, van Deemter and Kibble 2000, 632), as in (15) (ibid.), but I would argue

⁹This is not to say that there are not several datasets which mark up coreference for predicates (see below); however since most kinds of predication are not covered by OntoNotes, mainstream papers, systems and the CoNLL shared tasks have not included these cases. For examples of how I think it should be handled, see the Appendix, as well as Section 4.2

that it is more broadly a problem of “change of scope”, since modal scopes create the same type of issue, as in (16) (a made-up example).¹⁰

- (15) [Henry Higgins, who was formerly [sales director of Sudsy Soaps]_i], became [president of Dreamy Detergents]_i
- (16) If [Beyoncé]_i were [the Queen of England]_i, [she]_i would....

In (15), Henry Higgins is not simultaneously sales director at one company and president of another, raising doubts as to the value of a coreference chain including all of those NPs. Likewise in (16), the pronoun “she” arguably corefers more to the Queen of England (in a world where Beyoncé is the Queen), but not to the first mention of Beyoncé. In fact, this case verges on examples of non-referential bound anaphora, also discussed by van Deemter and Kibble, but which cannot be discussed here for space reasons.

So should we rule out predication due to such problems? After careful consideration of different opinions in the reviews of this paper, I believe the best answer is ‘yes, predication is different from other types of coreference’, but also ‘no, we cannot ignore it if we want to get the full picture’. My reluctance to ignore predication is based on a double dissociation: not all cases of predication raise this problem, and the problem can arise without predication. The most obvious case for including predication is in identificational predicates, which in English usually involves a definite predicate, but in other languages (e.g. Japanese) can often be indistinguishable from other predications. Compare these cases, marked up according to ON guidelines, which include naming predicates as copular (BBN Technologies, 2007, 27):

- (17) [Elizabeth II]_i is the Queen of England. [She]_i ...
- (18) The Queen of England is [Elizabeth II]_i. [She]_i ...
- (19) She was crowned [Elizabeth II]_i in 1953. [She]_i ...

Although they forbid annotating subject and predicate as a coreferent pair, ON guidelines do specify a hierarchy for determining which of the two should be taken as the antecedent for subsequent mentions, ranking names above pronouns, and pronouns above definite NPs. These kinds of configurations therefore appear in the official CoNLL coreference shared task dataset and consequently model the output that contemporary coreference resolution systems attempt to produce. I will leave it to readers to imagine scenarios in which the missing member of the chain will lead to loss of information,¹¹ but suffice it to say that including them would not lead to the scope problems above; in fact, some corpora already separate indefinite predication from regular coreference such as ARRAU (Poesio and Artstein, 2008) and GUM, which treat identificational predication as regular coreference, and either label predicative markables whenever semantically applicable (ARRAU), or maintain a special coreference relation subtype for non-identificational cases (GUM, see Appendix).

For the other side of the dissociation, we can consider chains of simple, definite NPs, in which the scope problem does arise, as in (20).

¹⁰Massimo Poesio has pointed out that the problem is even broader and ultimately stems from the multiple types of entity expressions in names or variables, truth value expressions and quantification in the sense of Montagovian type shifting, see Partee (1986). In this discussion I will limit myself to the prominent and frequent issues in narrowly defined copula predication, but indeed similar problems arise for negated NPs, relational indefinites and other ‘tricky NPs’ (see Landman 2004).

¹¹See also Lee et al. (2013, 902) for an example.

- (20) [A fresh major in the Swedish army], in 1812 [Gordon] went to war . . . In 1875 [the now general in the Russian army] was ready to pursue [his] ultimate achievement. . . [Gordon] is buried in. . . (adapted from GUM_bio_gordon)

We can introduce temporally inconsistent mentions (the buried man is not a Swedish major, the Swedish major is not a Russian general), and similarly modality inconsistencies or ambiguous mentions, etc. Although predication, naming, objects of ‘as’ and other constructions are *likely* locations for this in English, they are not the problem in themselves (so not a necessary condition), and they do not guarantee that there is a problem (so not sufficient). I agree such problems need to be addressed (see Section 4), but with few exceptions, the solution in most major datasets so far has been to largely ignore them, rather than to mark them up as a special case of anaphora (similarly to bridging, discourse deixis or split antecedents).

3. Isn’t this someone else’s job?

Before suggesting possible solutions, I would like to briefly outline why two potential objections to this paper’s proposals do not offer alternative solutions to the problems raised here. Although there may be other suggestions on how to deal with the phenomena not covered in OntoNotes, most probably follow one of two ways of ‘punting’ the issues: either to syntax or semantics.

3.1 Can’t we get all of this from syntax?

This line of reasoning is often raised in defence of coreference datasets which exclude predication, but which do have gold treebanking information: since syntax trees represent predication, naming constructions, nesting, and other structures more or less unambiguously, can’t we just leave them out of the coreference annotation proper and recover them from the trees?

The answer is no. I have yet to see a convincing case where this has been done, and there are good reasons why we should not think that it is possible. Contrast the following pairs, adapted from GUM:

- (21) a. [He] would be a Libertarian today (*no coref, since this is hypothetical: he is *not* a Libertarian*)
 b. [The principles governing an F-E translation]_i would then be: [reproduction of grammatical units; consistency in word usage; and meanings in terms of the source]_i (*coref, since in fact, those are the principles of F-E translation*)
- (22) a. [This coffee table] is glass (*not predicative coref, only specifies the substance that most/part of the table is made of*)
 b. [This ice here]_i of course is [water]_i (*predicative coref, part of a chemistry demonstration in which the speaker literally identifies an ice cube as being the same water in a solidified state*)
- (23) [He]_i was not the leaf-collecting doctor, but [an altogether strange man, with silver eyebrows in his smooth face and long fine-knuckled hands]_i (GUM_fiction_lunre)

In (21)–(22) the same syntactic construction yields coreferring expressions in one case and no coreference in the other (in (21), identity coreference, in (22), predication coreference, labeled as such in GUM). In (23), only part of a coordinate predicate NP is coreferring, meaning that extracting

the correct span using only syntax is non-trivial, and of course the phenomena in these examples can co-occur (imagine processing modal coordinate substance predications correctly!). Syntactic approaches also assume a well-formed syntax tree uttered contiguously by a single speaker, which in some cases is not a given.

It is also extremely difficult to know whether fairly mundane spatio-temporal predicate NPs are coreferring, regardless of definiteness:

- (24) This town is 35 minutes from the harbor (*no coref*, *purely spatio-temporal ‘is’, since town \neq 35 minutes*)
- (25) But Christmas is still the whole winter to wait (*no coref*, *but most syntax trees would show ‘winter’ as a definite predicate NP*)

Similarly for NP-internal or nested coreference resolution, we could easily have ambiguities in OntoNotes-style data. Consider this minor modification of the ‘sheep’ example above

- (26) A sheep_i in a coat_j with a star on its_{i/j} back

The data in (26) means that nested resolution is needed for disambiguation and cannot be extracted from syntax automatically, just as it is for predication. Finally for compound modifiers, although many coreferential cases include verbatim repetition of an entire compound, which can perhaps be recognized from word forms alone, some cases do not, with examples like (10) above forming the most striking cases. There can be no syntactic solution if we want to recover such relations.

To be clear, this is not to say that morpho-syntactic criteria can never be used in coreference annotation guidelines in any way: the reason why projects have used syntactic constructions diagnostically is because they can be very helpful in differentiating subtly different constructions and as a result help raise inter-annotator agreement. The point is rather that use of such criteria should be in service of semantic distinctions, which should be our actual object of interest: syntax and morphology should not be allowed to rule out things which we are certain do co-refer (for example morphologically incongruent singular and plural NPs), and they should not be used to admit things that do not (for example appositions which do not corefer). And if something is supposed to be recoverable from syntax, this should not move us to exclude it from coreference annotation – otherwise we are effectively mandating that any coreference resolution system will also have to tackle syntactic parsing.

3.2 How about semantics?

Another line of reasoning is that for ‘quirky’ coreference cases, including predication but also various kinds of distributive semantics (which space prevents me from discussing), semantic annotation should be in charge of annotating predicate structure in a way that disambiguates the issues in Section 3.1. The problem here is that even the most elaborate semantic annotation formalisms available right now do not address the main problems, since predication is largely subsumed in lexical semantics. In PropBank (Palmer et al., 2005), one of the simpler but most widespread types of semantic annotation, predicates are simply argument structure graphs with word sense disambiguation, but copula ‘be’ is simply annotated as BE.01, as in the following example from OntoNotes:

- (27) a. the most special is rice (ON, cctv_0000)
 b. Prop: BE.01, ARG1: ‘the most special’, ARG2: ‘rice’

It is therefore impossible to know whether this is a coreferential case or not (it is, specifically a non-identification predicative coreference: in this case a rice plant is being discussed).

More complex semantic analysis is undertaken in the less widespread but highly detailed formalism of Abstract Meaning Representation (AMR, Banarescu et al. 2013), which has facilities to indicate, for example, predication negation, possibly solving the problem in (23); but AMR is not aligned to words, and it is impossible to use it for the purpose of coreference resolution, as the following example from the Little Prince corpus (ibid.) illustrates:

- (28) a. It is my fault that you have not known it all the while.
 b. (f / fault-01
 :ARG1 (i3 / i)
 :ARG2 (k / know-01 :polarity -
 :ARG0 (y / you)
 :ARG1 (i2 / it)
 :time (w / while-away-01
 :duration (a / all))))

The AMR analysis captures the argument structure of ‘fault’ perfectly (arguments: the speaker and the ‘knowing’ predicate), but leaves no indication of the original expletive subject ‘it’, its relationship with the extraposed clause ‘that you have not known...’, etc. Because of AMR’s (intentional) distance from word forms in the original sentence, it does not fulfill or replace the task of text-anchored coreference resolution – AMR simply has no obligation to include nodes for each surface pronoun or NP. Similarly, neither of these formalisms analyzes the properties of discourse referents in terms of genericity, or contemporaneous temporal or modal scope.

4. Fixing the problems

Is it impossible to envision a semantic or syntactic formalism that can represent the facts discussed here? Certainly not: we could just extend semantic or syntactic annotation to distinguish such cases. But here I would like to argue that covering these cases is part of the job of coreference resolution, and realistically, if coreference resolution does not tackle them, they are unlikely to be covered by other annotation efforts in the current landscape.

4.1 Unexiling tricky coreference

Although it should be clear by now that I would like to see a lot more things annotated in coreference resolution datasets, I would like to emphasize that I am not saying that all of these types of coreference are *the same*. Addressing them does not necessarily mean that we need to lose the distinction between predication, pronominal anaphora and other types of coreference. OntoNotes itself distinguishes several types of apposition from identity coreference, and several other corpora distinguish not only appositions, but also pronominal anaphora from lexical identity (e.g. GUM) or non-identificational predication (e.g. both GUM and ARRAU for English, or CESS-ECE for Spanish, Recasens et al. 2007, which are all useful examples of explicitly handling, rather than ignoring predication), in addition to harder phenomena such as discourse deixis or bridging anaphora.

Critics may point out that including trickier cases will inevitably lead to lower agreement, but I would answer that doing so should not decrease agreement on ‘easy cases’, and that in instances

expression type	instances per 1K tokens	% of total
<i>pron. anaphora</i>	59.84	44.98%
<i>cataphora</i>	1.39	1.05%
<i>nominal predicate</i>	4.94	3.71%
<i>compound mod.</i>	5.79	4.36%
<i>split antecedent</i>	1.22	0.92%
<i>apposition</i>	3.81	2.86%
<i>coref. name</i>	19.84	14.92%
<i>other indef. NP</i>	14.53	10.92%
<i>other coref</i>	21.67	16.29%
total	133.02	100%

Table 1: Coreference link type distribution in GUM

of disagreement, most often each of the conflicting analyses corresponds to a valid reading of an ambiguous context (as demonstrated by Poesio and Artstein 2005). Human annotators also tend to disregard or not notice a relation more often than connecting non-corefering entities (Jiang et al., 2013; Zeldes, 2017), meaning data would still have relatively high precision (for GUM, which includes the phenomena discussed here, student annotators achieved precision, recall and F1 of 0.918/0.811/0.861 compared to the adjudicated gold standard, when minimal mention span matching is used, Zeldes 2017, 596). With this in mind, including these phenomena and labeling them as such, as is done in ARRAU or GUM, seems in essence already close to the ideal solution, with a few caveats and additions.

The first is that the default coreference resolution task should, in my opinion, cover all types of cluster-based¹² coreference discussed here, including compound modifiers, indefinites and predication. As shown in Table 1, which tallies types of anaphoric expressions in GUM, all of these are extremely frequent, non-marginal phenomena, which often overlap with clear cases of identity coreference (e.g. definite copula predicates, semantically specific compound modifiers). Taken together, indefinite NPs, compound modifiers and predicates form around 19% of all previously mentioned NPs, occurring about 26 times per 1K tokens. For comparison, this is more than the remaining definite common NPs (21.67 times) or subsequently mentioned proper names (19.84). Nested dates (not shown in the table) are also ubiquitous in GUM data, with 13.8% of temporal NPs being nested in other temporal NPs (including singletons) and 38% of corefering year expressions in the corpus being nested in a date entity span, which would exclude them in ON. Removing such mentions from the target standard for mainstream NLP leads to counterintuitive results and ultimately creates an incongruity between systems’ performance in shared tasks and their coverage for real world applications, which often assume the premise of being able to collect *all* information about an entity in a text (say, some year) when coreference resolution is ‘correct’.

This position, if accepted, raises the question of what we should do about Henry Higgins, president of Dreamy Detergents? Haven’t we been here before, and didn’t the criticism of the ACE dataset teach us that ‘change over time’ makes text-anchored coreference futile?

¹²I use this term to set aside other types of anaphora, such as bridging, sense anaphora (e.g. ‘one’ anaphora), etc., not because they are not important, but because we already have much work to do on many types of identity coreference and closely related concepts.

4.2 Addressing Scope

The ‘change over time’, or better, ‘change of scope’ problem raised by van Deemter and Kibble (2000) remains the biggest challenge to a more inclusive view of coreference, and I have answers for it on several levels. On the practical level, scope problems are both not very common and ubiquitous. They are not very common since true hard cases, such as texts covering an entity changing substantially vis-a-vis the relevance of referring expressions over time, are only a subset of texts, and even in those texts, any problematic entity often co-exists with scores of other predicate NPs which do not raise such problems. Refusing to include those hundreds of identificational predicate NPs is throwing out the baby with the bath water.

At the same time, the problem is ubiquitous (cf. Recasens et al. 2012 on ‘near-identity coreference’). Any long text will involve changes to participants, which, even if they are not stated, mean that not all predicates stated of some entity actually apply simultaneously. If we omit predicates from the coreference annotation and then tell the story of Henry Higgins, he does not remain the same Henry Higgins whether we include the information about his various jobs or not, or whether we express that information using NPs or not – we could just say he was ‘fired’ and alter the properties of the entity (i.e. his job) without interacting with coreference. To then say that we will use syntax or semantics to describe Henry’s biography does not absolve us of the need to address scope, since the same inconsistencies will result in a semantic annotation of the text. But ignoring ‘difficult’ predicate NPs as part of coreference annotation will lead to many gaps in what practitioners using coreference resolution might expect or could harness.

Moving over to the theoretical perspective, I think there is enough merit in representing scope that we as a community of researchers interested in coreference should consider what is the right thing to do. The fact that we recognize the problem with Henry Higgins suggests that we *do* know that both the ‘sales director’ Henry and the ‘president’ Henry are the same person: just not at the same time. It could be tempting to simply allow a co-indexed scope attribute in our coreference annotation schemes, as in (29) for the Beyoncé example.

- (29) If `<coref id="1">Beyoncé</coref>` were `<coref id="1" scope="s1">the Queen of England</coref>`, `<coref id="1" scope="s1">she</coref>` would conduct the annual swan upping.

This notation would accurately show that the pronoun ‘she’ refers to its intuitive antecedent entity in the document context, ‘Beyoncé’, while also indicating that this specific ‘she’ is referential in the world in which Beyoncé is the Queen, and co-refers to that Queen. We could also add a coreference type indicating that we are looking at identity-predication, etc.

For modal cases, adding a scope ID may be enough, but for temporal ones, we could even consider annotating intervals where these are known (or place-holders where they are not; see the RED corpus for a similar strategy in event annotation, O’Gorman et al. 2016):

- (30) `<coref id="1">Henry Higgings, who was formerly <coref id="1" scope="s1" from="1999" to="2001">sales director at Sudsy Sopes</coref></coref>`, became `<coref id="1">President of Dreamy Detergents</coref>`

Omitting scope IDs and intervals would mean that these are the ‘general’ scope conditions of the document, and that further information is not known.

Alternatively, given that there have been independent efforts to annotate modal (Hendrickx et al., 2012; Nissim et al., 2013; Rubinstein et al., 2013) and temporal scope (Pustejovsky and Stubbs, 2011; Styler IV et al., 2014; Pustejovsky et al., 2019) for semantics, it may make sense to encode scope independently of coreference annotation, as a span level property of sentences, or VPs or any text span, which interacts with coreference, since many other types of annotation could be affected by modal and temporal scopes. For example, the hypothetical ‘conduct the swan upping’ (an annual census of the Queen’s swans) in (29) is also restricted to the modal scope in which Beyoncé is the Queen, and event annotations of this predicate could take this into account.

I do not want to pretend that annotating scope for coreference is simple, or that we are in a position to easily add it to large benchmark datasets such as OntoNotes: there are surely many complications, even without which this would mean much work, and a detailed examination of how coreference interacts with existing scope annotation schemes is certainly in order. But assuming we can take the position that (serious) scope problems are fairly rare, then I think the right starting point before moving on to an adequate representation for these cases is to include the phenomena missing in OntoNotes as an ideal target, and offer scope as a topic for further advanced research on coreference, similarly to bridging anaphora or split antecedents, whose absence in benchmarks like OntoNotes is well understood, but does not interfere with progress on the most common types of coreferentiality in corpora.

5. Conclusion

If this opinion piece falls short of changing any coreference annotation practices, then I hope it at least serves one purpose: to make researchers aware of the limitations of ON-style coreference, and by extension most NLP tools for coreference resolution. My experience has been that these limitations are not well understood by proficient NLP practitioners who use tools such as vanilla e2e coref (Lee et al., 2017), AllenNLP (Gardner et al., 2018) or Spacy’s Huggingface (Clark and Manning, 2016) implementations, and this often leads to surprises. Conscious choices for or against including some of the trickier phenomena will inevitably involve factors relating to trade-offs of linguistic adequacy versus reliability and application domains, but informed decisions are better than ones made spontaneously or based on inertia.

A more ambitious goal for this paper is that readers involved in the production and (re)annotation of coreference datasets will consider following corpora such as GUM and ARRAU in taking a broader view of the phenomenon, and use subtypes to carve out areas which some systems might want to leave out or label in special ways for downstream applications – the Appendix offers some worked examples from the GUM corpus as a starting point. Including more coreference phenomena does not mean that we have to be naive about the potentially lower level of agreement on complex cases (e.g. generics, predication); but removing such cases from output that contains them is generally much easier than trying to add them to data in which they are not included in any way, and if modifying data along different lines benefits specific subsets of downstream tasks, then the flexibility offered by distinguishing but including tricky phenomena is all the more valuable. This has also been one of the main lessons I have learned from working on multilayer corpora (see Zeldes 2018): no, we cannot get everything from syntax as pointed out above, but syntax is great for finding specific phenomena or manipulating data into different schemes (e.g. removing compound modifiers, changing how appositions are handled). In one paper, this has even allowed transforming GUM,

the dataset most closely matching this proposal, to follow the OntoNotes guidelines (the resulting dataset is called ‘OntoGUM’, see Zhu et al. 2021 for details and evaluation).

Finally, I think that as corpus resources converge across languages on uniform standards, as evidenced in the Universal Dependencies project for treebanking, POS tagging and morphological annotation (de Marneffe et al., 2021), or multilingual initiatives in other areas of discourse (e.g. the DISRPT shared tasks on discourse relations, Zeldes et al. 2019, 2021) we will need to make coreference more consistent across datasets. This will mean that we have to minimize the use of guidelines based on language-specific forms, such as definiteness or specific constructions. Instead, I have argued here that coreference should primarily be decided on semantic grounds, which seem likelier to transfer between languages. As an added motivation, I offer the consideration that cross-linguistically and semantically grounded coreference annotations are likely to work better for applications involving multiple corpora, such as multilingual and cross-document coreference, as well as entity linking (see Lin and Zeldes 2021 for more on merging entity linking with cross-document coreference).

The problems discussed in this paper are in my opinion not intrinsically ones of copula sentences or modal verbs or other constructions, but rather issues at the discourse-semantics interface, where meanings from individual sentences coalesce to form narratives in which entities change. This suggests that tackling the issues in a cross-linguistically applicable way will ultimately require the development of standards for handling spatiotemporal, modal and other types of scope independently of language or text type.

References

- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. Abstract Meaning Representation for Sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria, 2013.
- BBN Technologies. OntoNotes English co-reference guidelines version 7.0. Technical report, BBN Technologies, 2007.
- BBN Technologies. OntoNotes Chinese co-reference guidelines version 4.0. Technical report, BBN Technologies, 2008.
- Archana Bhatia, Chu-Cheng Lin, Nathan Schneider, Yulia Tsvetkov, Fatima Talib Al-Raisi, Laleh Roostapour, Jordan Bender, Abhimanu Kumar, Lori Levin, Mandy Simons, and Chris Dyer. Automatic classification of communicative functions of definiteness. In *Proceedings of COLING 2014*, pages 1059–1070, Dublin, Ireland, 2014.
- Kevin Clark and Christopher D. Manning. Deep reinforcement learning for mention-ranking coreference models. In *Proceedings of EMNLP 2016*, pages 2256–2262, Austin, TX, 2016.
- Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. Universal Dependencies. *Computational Linguistics*, 47(2):255–308, 2021.
- George R. Doddington, Alexis Mitchell, Mark A. Przybocki, Lance A. Ramshaw, Stephanie Strassel, and Ralph M. Weischedel. The Automatic Content Extraction (ACE) program. tasks,

- data, and evaluation. In *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC-2004)*, pages 837–840, 2004.
- John W. Du Bois, Wallace L. Chafe, Charles Meyer, Sandra A. Thompson, Robert Englebretson, and Nii Martey. Santa barbara corpus of spoken american english, parts 1-4, 2000-2005.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. AllenNLP: A deep semantic natural language processing platform. In *Proceedings of ACL 2018*, pages 1–6, Melbourne, 2018.
- Iris Hendrickx, Amália Mendes, and Silvia Mencarelli. Modality in text: a proposal for corpus annotation. In *Proceedings of LREC 2012*, pages 1805–1812, Istanbul, Turkey, 2012.
- Lynette Hirschman and Nancy Chinchor. Appendix F: MUC-7 coreference task definition (version 3.0). In *Seventh Message Understanding Conference (MUC-7)*, 1998.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. OntoNotes: The 90% solution. In *Proceedings of NAACL 2006: Short Papers*, pages 57–60, New York, 2006.
- Lili Jiang, Yafang Wang, Johannes Hoffart, and Gerhard Weikum. Crowdsourced entity markup. In *Proceedings of Crowdsourcing the Semantic Web*, pages 59–68, Sydney, 2013.
- Mandar Joshi, Omer Levy, Luke Zettlemoyer, and Daniel Weld. BERT for coreference resolution: Baselines and analysis. In *Proceedings of EMNLP-IJCNLP 2019*, pages 5803–5808, Hong Kong, China, 2019.
- Manfred Krifka. Basic notions of information structure. *Acta Linguistica Hungarica*, 55:243–276, 2008.
- Kerstin Kunz. *Variation in English and German Nominal Coreference. A Study of Political Essays*. Peter Lang, Frankfurt, 2009.
- Fred Landman. *Indefinites and the Type of Sets*. Blackwell, Malden, MA, 2004.
- Heeyoung Lee, Angel Chang, Yves Peirsman, Nathanael Chambers, Mihai Surdeanu, and Dan Jurafsky. Deterministic coreference resolution based on entity-centric, precision-ranked rules. *Computational Linguistics*, 39(4):885–916, 2013.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. End-to-end neural coreference resolution. In *Proceedings of EMNLP 2017*, pages 188–197, Copenhagen, 2017.
- Jessica Lin and Amir Zeldes. WikiGUM: Exhaustive entity linking for Wikification in 12 genres. In *Proceedings of The Joint 15th Linguistic Annotation Workshop (LAW) and 3rd Designing Meaning Representations (DMR) Workshop*, pages 170–175, Punta Cana, Dominican Republic, 2021.
- Nafise Sadat Moosavi and Michael Strube. Which coreference evaluation metric do you trust? A proposal for a link-based entity aware metric. In *Proceedings of ACL 2016*, pages 632–642, Berlin, 2016.

- Anna Nedoluzhko, Jiří Mírovský, Radek Ocelák, and Jiří Pergler. Extended coreferential relations and bridging anaphora in the Prague Dependency Treebank. In *Proceedings of the 7th Discourse Anaphora and Anaphor Resolution Colloquium (DAARC 2009)*, pages 1–16, Goa, India, 2009.
- Anna Nedoluzhko, Michal Novák, Martin Popel, Zdeněk Žabokrtský, Amir Zeldes, and Daniel Zeman. CorefUD 1.0: Coreference meets Universal Dependencies. In *Proceedings of LREC 2022*, Marseille, France, 2022.
- Azadeh Nikfarjam, Ehsan Emadzadeh, and Graciela Gonzalez. Towards generating a patient’s timeline: Extracting temporal relationships from clinical notes. *Biomedical Informatics*, 46:40–47, 2013.
- Malvina Nissim, Paola Pietrandrea, Andrea Sansò, and Caterina Mauri. Cross-linguistic annotation of modality: a data-driven hierarchical model. In *Proceedings of the 9th Joint ISO - ACL SIGSEM Workshop on Interoperable Semantic Annotation*, pages 7–14, Potsdam, Germany, 2013.
- Michal Novák and Anna Nedoluzhko. Correspondences between Czech and English coreferential expressions. *Discours*, 16:3–41, 2015.
- Jan-Ola Östman. *You Know: A discourse-functional approach*. Pragmatics & Beyond, II:7. John Benjamins, Amsterdam, 1981.
- Tim O’Gorman, Kristin Wright-Bettner, and Martha Palmer. Richer event description: Integrating event coreference with temporal, causal and bridging annotation. In *Proceedings of the 2nd Workshop on Computing News Storylines (CNS 2016)*, pages 47–56, Austin, TX, 2016.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. The Proposition Bank: A corpus annotated with semantic roles. *Computational Linguistics*, 31(1):71–106, 2005.
- Barbara H. Partee. Noun phrase interpretation and type-shifting principles. In Jeroen Groenendijk, Martin Stokhof, and Dick de Jongh, editors, *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*, pages 115–143. Foris, Dordrecht, 1986.
- Massimo Poesio and Ron Artstein. The reliability of anaphoric annotation, reconsidered: Taking ambiguity into account. In *Proceedings of the Workshop on Frontiers in Corpus Annotation II: Pie in the Sky*, pages 76–83, Stroudsburg, PA, 2005. ACL.
- Massimo Poesio and Ron Artstein. Anaphoric annotation in the ARRAU corpus. In Nicoletta Calzolari, Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odjik, Stelios Piperidis, and Daniel Tapias, editors, *Proceedings of LREC 2008*, pages 1170–1174, Marrakesh, Morocco, 2008.
- Massimo Poesio, Sameer Pradhan, Marta Recasens, Kepa Rodriguez, and Yannick Versley. Annotated corpora and annotation tools. In Massimo Poesio, Roland Stuckardt, and Yannick Versley, editors, *Anaphora Resolution: Algorithms, Resources, and Applications*, Theory and Applications of Natural Language Processing, pages 97–140. Springer, Heidelberg, 2016.
- Paul M. Postal. Anaphoric islands. In Robert I. Binnick, editor, *5th regional meeting of the Chicago Linguistic Society*, volume 5, pages 205–239. Chicago Linguistic Society, Chicago, 1969.

- Sameer Pradhan, Lance Ramshaw, Mitchell Marcus, Martha Palmer, Ralph Weischedel, and Nianwen Xue. CoNLL-2011 shared task: Modeling unrestricted coreference in OntoNotes. In *Proceedings of the Conference on Computational Natural Language Learning: Shared Task*, pages 1–27, Portland, OR, 2011.
- Sameer Pradhan, Xiaoqiang Luo, Marta Recasens, Eduard Hovy, Vincent Ng, and Michael Strube. Scoring coreference partitions of predicted mentions: A reference implementation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 30–35, Baltimore, MD, 2014.
- James Pustejovsky and Amber Stubbs. Increasing informativeness in temporal annotation. In *5th Linguistic Annotation Workshop (LAW V)*, pages 152–160, Portland, OR, 2011.
- James Pustejovsky, Ken Lai, and Nianwen Xue. Modeling quantification and scope in abstract meaning representations. In *Proceedings of the First International Workshop on Designing Meaning Representations (DMR)*, pages 28–33, Florence, Italy, 2019.
- Swayambhu Nath Ray, Shib Sankar Dasgupta, and Partha Talukdar. AD3: Attentive deep document dater. In *Proceedings of EMNLP 2018*, pages 1871–1880, Brussels, Belgium, 2018.
- Marta Recasens, M. Antònia Martí, and Mariona Taulé. Where anaphora and coreference meet. annotation in the Spanish CESS-ECE corpus. In *Proc. RANLP 2007*, Borovets, Bulgaria, 2007.
- Marta Recasens, M. Antònia Martí, and Constantin Orasan. Annotating near-identity from coreference disagreements. In *Proceedings of LREC 2012*, pages 165–172, Istanbul, Turkey, 2012.
- Tanya Reinhart. Pragmatics and linguistics. An analysis of sentence topics. *Philosophica*, 27:53–94, 1982.
- Aynat Rubinstein, Hillary Harner, Elizabeth Krawczyk, Daniel Simonson, Graham Katz, and Paul Portner. Toward fine-grained annotation of modality in text. In *Proceedings of the Workshop on Annotation of Modal Meanings in Natural Language*, pages 38–46, Potsdam, Germany, 2013.
- Veselin Stoyanov, Nathan Gilbert, Claire Cardie, and Ellen Riloff. Conundrums in noun phrase coreference resolution: Making sense of the state-of-the-art. In *Proceedings of ACL/IJCNLP 2009*, pages 656–664, Suntec, Singapore, 2009.
- William F. Styler IV, Steven Bethard, Sean Finan, Martha Palmer, Sameer Pradhan, Piet C de Groen, Brad Erickson, Timothy Miller, Chen Lin, Guergana Savova, and James Pustejovsky. Temporal annotation in the clinical domain. *TACL*, 2:143–154, 2014.
- Kees van Deemter and Rodger Kibble. On coreferring: Coreference in MUC and related annotation schemes. *Computational Linguistics*, 26(4):629–637, 2000.
- Gregory Ward, Richard Sproat, and Gail McKoon. A pragmatic analysis of so-called anaphoric islands. *Language*, 63(3):439–474, 1993.
- Ralph Weischedel, Sameer Pradhan, Lance Ramshaw, Jeff Kaufman, Michelle Franchini, Mohammed El-Bachouti, Nianwen Xue, Martha Palmer, Jena D. Hwang, Claire Bonial, Jinho Choi,

- Aous Mansouri, Maha Foster, Abdel aati Hawwary, Mitchell Marcus, Ann Taylor, Craig Greenberg, Eduard Hovy, Robert Belvin, and Ann Houston. Ontonotes release 5.0. Technical report, Linguistic Data Consortium, Philadelphia, 2012.
- Wei Wu, Fei Wang, Arianna Yuan, Fei Wu, and Jiwei Li. CorefQA: Coreference resolution as query-based span prediction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6953–6963, Online, 2020.
- Annie Zaenen. Mark-up barking up the wrong tree. *Computational Linguistics*, 32(4):577–580, 2006.
- Amir Zeldes. The GUM corpus: Creating multilayer resources in the classroom. *Language Resources and Evaluation*, 51(3):581–612, 2017.
- Amir Zeldes. *Multilayer Corpus Studies*. Routledge Advances in Corpus Linguistics 22. Routledge, London, 2018.
- Amir Zeldes and Shuo Zhang. When annotation schemes change rules help: A configurable approach to coreference resolution beyond OntoNotes. In *Proceedings of the Workshop on Coreference Resolution Beyond OntoNotes (CORBON)*, pages 92–101, San Diego, CA, 2016.
- Amir Zeldes, Debopam Das, Erick Galani Maziero, Juliano Desiderato Antonio, and Mikel Iruskieta. The DISRPT 2019 Shared Task on elementary discourse unit segmentation and connective detection. In *Proceedings of Discourse Relation Treebanking and Parsing (DISRPT 2019)*, pages 97–104, Minneapolis, MN, 2019.
- Amir Zeldes, Yang Janet Liu, Mikel Iruskieta, Philippe Muller, Chloé Braud, and Sonia Badene, editors. *Proceedings of the 2nd Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2021)*, Punta Cana, Dominican Republic, 2021.
- Rui Zhang, Cícero Nogueira dos Santos, Michihiro Yasunaga, Bing Xiang, and Dragomir Radev. Neural coreference resolution with deep biaffine attention by joint mention detection and mention clustering. In *Proceedings of ACL 2018*, pages 102–107, Melbourne, Australia, 2018.
- Shuo Zhang and Amir Zeldes. GitDOX: A linked version controlled online XML editor for manuscript transcription. In *Proceedings of FLAIRS-30*, pages 619–623, Marco Island, FL, 2017.
- Yilun Zhu, Sameer Pradhan, and Amir Zeldes. Anatomy of OntoGUM—Adapting GUM to the OntoNotes scheme to evaluate robustness of SOTA coreference algorithms. In *Proceedings of the Fourth Workshop on Computational Models of Reference, Anaphora and Coreference (CRAC 2021)*, pages 141–149, Punta Cana, Dominican Republic, 2021.

Appendix A. Worked examples

This section gives detailed examples taken from different genres in the GUM corpus, which includes two versions of coreference annotations: GUM’s native, exhaustive coreference scheme, which follows the suggestions in this paper, and an automatically produced version of the same data following the OntoNotes guidelines, a rendition of the corpus referred to as OntoGUM (see Zhu et al. 2021 for details). Markup for all examples in this appendix can be found at the GUM corpus repository, in both the original and OntoNotes scheme, at <https://github.com/amir-zeldes/gum>.

A.1 How to Grow Basil

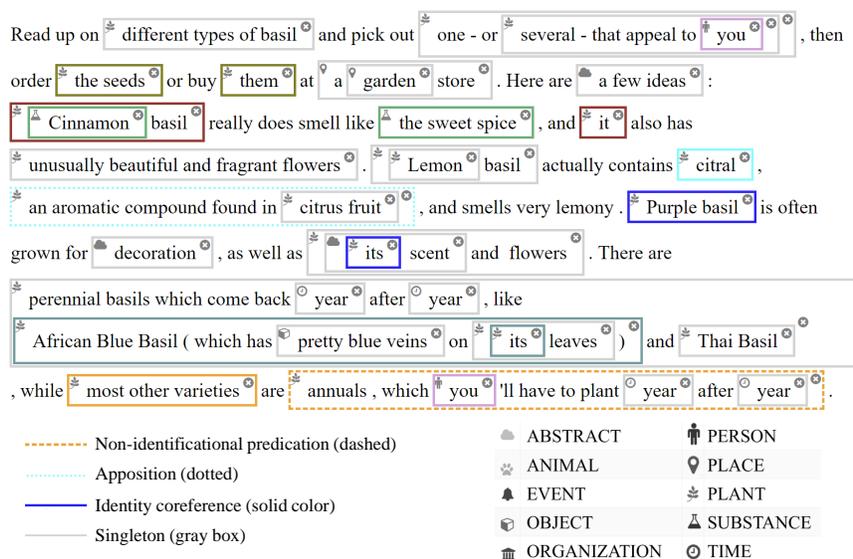


Figure 1: GUM-style coreference annotations (GUM_who_basil), visualized using GitDox (Zhang and Zeldes, 2017).

This excerpt from a Wikihow guide on how to grow basil (see the corpus website for source and licensing information on all texts) illustrates the proposal’s coverage for nested identity coreference and non-identificational predicative coreference, as well as singletons and entity type annotations found in GUM. The more sparse OntoGUM representation illustrates the comparative loss of information based on the OntoNotes scheme.

Beyond the singletons (in grey) and entity types (icons in top left corners of boxes) found in Figure 1 (see the legend), some items to note include the common noun compound modifier antecedent ‘Cinammon’, the predication between ‘most other varieties’ and ‘annuals...’, and the generic pronoun ‘you’ (in pink). Additionally, we see nested coreference in ‘[African Blue Basil (which has pretty blue veins on [its]_i leaves)]_i’. In Figure 2 we see that the nested coreference set is not included according to the OntoNotes scheme, since it is not mentioned a subsequent time outside the nesting mention, and is therefore considered a singleton (and singletons are generally omitted). Similarly the mention of ‘Cinammon’ is ignored, and predication for ‘annuals’ cannot be recovered trivially (see Section 2.4).

Read up on different types of basil and pick out one - or several - that appeal to you , then order the seeds or buy them at a garden store . Here are a few ideas : Cinnamon basil really does smell like the sweet spice , and it also has unusually beautiful and fragrant flowers . Lemon basil actually contains citral , an aromatic compound found in citrus fruit , and smells very lemony . Purple basil is often grown for decoration , as well as its scent and flowers . There are perennial basil which come back year after year , like African Blue Basil (which has pretty blue veins on its leaves) and Thai Basil , while most other varieties are annuals , which you 'll have to plant year after year .

Figure 2: OntoNotes-style coreference annotations (GUM_who_basil).

A.2 Spoken indefinites

This example comes from a conversation (GUM_conversation_lambda), originally from the UCSB corpus of Spoken American English (Du Bois et al., 2000-2005), and included in GUM’s conversation genre.

- A. *Miles:* Anyway, * this guy, I mean he comes over there, immediately afterwards —
- Jamie:* The same guy who got his — the same guy who's —
- Miles:* No, it's just some random guy.
- Harold:* Or just some random guy.
- Miles:* No, a random guy, who was like sitting about * where you are. Yeah, He was sitting there
-
- B. *Miles:* Anyway, this guy, I mean he comes over there, immediately afterwards —
- Jamie:* The same guy who got his — the same guy who's —
- Miles:* No, it's just some random guy.
- Harold:* Or just some random guy.
- Miles:* No, a random guy, who was like sitting about where you are. Yeah, He was sitting there

Figure 3: Analyses of a conversation fragment according to coreference schemes based on GUM (above) and OntoNotes (below). Referring expressions with a gold star are annotated as accessible in context.

In Figure 3A, mentions of the ‘random guy’ are all grouped together, regardless of definiteness or predication, since these are not criteria for coreference in the GUM scheme, which follows the proposal in this paper. The OntoNotes version ignores some mentions (if they are indefinites in chain-medial position) and oddly splits the chain into two sets, creating the impression that there are two random guys, in addition to the person being asked about by the second speaker (there are only two guys in context). GUM’s mention annotation scheme, which marks up information status (with six subtypes), additionally flags referring expressions as accessible in context, including an initial deictic realization of ‘this guy’ and the situationally accessible ‘where you are’, which are marked by gold stars in the figure.

A.3 Datelines and headlines

This example comes from the news genre (taken from Wikinews) and illustrates two common problems in newswire data.

In Figure 4, ‘2015’ corefers with ‘this year’, and ‘this process’ corefers with both spans ‘a public process’ and the ‘process to consider...’ which is part of the headline. Because the headline and first subsequent mention of the ‘process’ are both indefinite, the first mention is eliminated in the OntoNotes version in Figure 5. Additionally, because the year in the dateline is nested in a date, OntoNotes guidelines rule it out as an antecedent for ‘this year’, meaning that the latter becomes a singleton, and both mentions are excluded.

New Zealand begins process to consider changing national flag design Thursday, May 7, 2015. On Tuesday, the New Zealand government announced the start of a public process to suggest designs for a new national flag, and determine whether their citizens would prefer a different national flag over the current one. The current flag of New Zealand. The current New Zealand flag is partially based on the United Kingdom's flag; the new one would be unique to New Zealand. The government's Flag Consideration Project has planned a number of conferences and roadshows as part of this process ... this year.

Figure 4: GUM-style coreference annotations (GUM_news_flag).

New Zealand begins process to consider changing national flag design Thursday, May 7, 2015. On Tuesday, the New Zealand government announced the start of a public process to suggest designs for a new national flag, and determine whether their citizens would prefer a different national flag over the current one. The current flag of New Zealand. The current New Zealand flag is partially based on the United Kingdom's flag; the new one would be unique to New Zealand. The government's Flag Consideration Project has planned a number of conferences and roadshows as part of this process ... this year.

Figure 5: OntoNotes-style coreference annotations (GUM_news_flag).