



HAL
open science

How to Evaluate Coreference in Literary Texts?

Ana-Isabel Duron-Tejedor, Pascal Amsili, Thierry Poibeau

► **To cite this version:**

Ana-Isabel Duron-Tejedor, Pascal Amsili, Thierry Poibeau. How to Evaluate Coreference in Literary Texts?. Computational Humanities Research, Dec 2023, Paris, France. hal-04367628

HAL Id: hal-04367628

<https://hal.science/hal-04367628>

Submitted on 30 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

How to Evaluate Coreference in Literary Texts?

Ana-Isabel Duron-Tejedor, Pascal Amsili and Thierry Poibeau

Lattice, CNRS & ENS-PSL & Université Sorbonne Nouvelle

Abstract

In this short paper, we examine the main metrics used to evaluate textual coreference and we detail some of their limitations. We show that a unique score cannot represent the full complexity of the problem at stake, and is thus uninformative, or even misleading. We propose a new way of evaluating coreference, taking into account the context (in our case, the analysis of fictions, esp. novels). More specifically, we propose to distinguish long coreference chains (corresponding to main characters), from short ones (corresponding to secondary characters), and singletons (isolated elements). This way, we hope to get more interpretable and thus more informative results through evaluation.

Keywords

Coreference, Evaluation, Metrics, Literature analysis

1. Introduction

Coreference resolution is a natural language processing (NLP) task that consists in determining the linguistic expressions, also called mentions, that refer to a same entity (Poesio et al., 2023). This task is crucial for NLP, and constitutes a logical complement to named entity recognition. For example, if one wants to track all the actions involving a given entity in a text, an accurate analysis of coreference is required, as most of the references to entities are usually done through coreference (i.e., pronouns and other linguistic items referring to an entity).

This is in particular the case for the analysis of fictions, especially novels: a proper coreference analysis is mandatory to analyze the plot, the evolution, and the interactions between the different characters (Bamman et al., 2014). Characters play a prominent roles in fiction. For this reason, we will only consider reference to characters in this study (which means we exclude, e.g., locations and company names, that are of course useful, but do not play a so prominent role in fictions).

The importance of coreference in NLP has been put forward by the Message Understanding Conferences (MUC) in the late 90ies (Grishman and Sundheim, 1996). Since then, numerous coreference resolvers have been developed for a variety of languages and coreference is a very active domain of research in NLP. A related issue is how to evaluate coreference. As coreference involves sets of mentions of various size linked together (i.e., referring to a same entity), traditional indicators like precision and recall have to be redefined, but there is no obvious way of doing it (should entities be weighted based on their relative importance in the text? What to do with singletons?). This led to several measures that do not always agree (i.e. one scoring method may indicate a good precision, whereas another method will not be so positive).



This has been already been demonstrated in various studies on the topic (Borovikova et al., 2022; Lion-Bouton et al., 2020; Moosavi and Strube, 2016; Moosavi, 2020). Here we want to stress complementary, albeit different, issues, and more specifically two of them:

- Existing systems and evaluation methods are based on small texts, and are not fully relevant when addressing long texts like novels;
- Existing evaluation methods tend to consider all the entities on an equal basis, regardless of the linguistic properties of the entity, which is not the case for most practical applications.

We will also point out a more general issue: the fact that none of the existing evaluation methods is really interpretable by a human.

The paper is structured as follows: we first remind the reader with state-of-the-art metrics and the main issues that have been reported about them. We then show that these measures have been defined for small texts but long texts like fictions tend to exacerbate the issues detailed in the previous section. We finish with some proposals based on a typology of coreference chains, to get more readable and actionable results.

2. State of the Art in Coreference Evaluation Methods

Coreference evaluation metrics can be divided in two categories: link-based and mention-based metrics. Roughly, link-based metrics represent entities as a set of links and mention-based represent them as a set of mentions. In this section, we just give a brief overview of the most popular methods and of their relative weaknesses, as more detailed overviews already exist. The interested reader can refer to (Borovikova et al., 2022; Lion-Bouton et al., 2020; Moosavi and Strube, 2016; Moosavi, 2020). Traditionally, in the literature on coreference, ‘key’ refers to the reference (gold annotations), and ‘response’ corresponds to the computed links.

MUC (Vilain et al., 1995) is only based on the minimum number of missing/extra links in the response compared to the key entities. It does not take into account singletons, nor the relative weight of the different entities (all the links have the same weight).

BLANC (Recasens and Hovy, 2011) is a metric that was built in a first place to evaluate coreference on key mentions. It was then adapted (Luo et al., 2014) to also evaluate response mentions. Contrary to *MUC*, *BLANC* uses non-coreferent links. The major problem with *BLANC* is that, when the number of gold mentions in the response entities is larger, the number of non-coreference links is also larger. In this case *BLANC* gives higher scores (for both recall and precision) without knowing if the gold mentions were resolved (Moosavi and Strube, 2016).

Two other measures, B^3 and *CEAF*, are mention-based. In the case of B^3 (Bagga and Baldwin, 1998), the overall recall/precision is computed based on the recall/precision of the individual mentions. For each mention m in the key entities, B^3 recall considers the fraction of the correct mentions that are included in the response entity of m . Similarly to *MUC*, B^3 precision is computed by switching the role of the key and response entities (Moosavi, 2020).

CEAF (Luo, 2005) relies on a one-to-one mapping, where each key entity is mapped to only one response entity ($CEAF_e$ is based on the mapping at the entity level, a variant $CEAF_m$ works at the mention level). Via this mapping, *CEAF* uses a similarity metric $\phi(i,j)$ to measure the similarity of the mapped entities with the reference. The one-to-one mapping at the heart of

CEAF leads to ignoring other correct decisions of unaligned response entities. Finally, $CEAF_e$ does not assign weights to the different entities: it can be seen as a metric that gives a global score of the proportion of entities that were found in the response (for recall) and the proportion of correct entities in the response (for precision).

All these metrics have been studied, compared and criticized, esp. in (Moosavi and Strube, 2016; Moosavi, 2020). They suffer from a number of known issues, especially what (Moosavi and Strube, 2016) call the “mention identification effect”. The authors propose the following experiment: they remove spurious mentions from the responses entities (the ones calculated automatically). This should automatically increase precision (since there are then less spurious mentions), while having no effect on recall. Surprisingly, the authors observe the following, “(i) recall changes for all the metrics except for MUC; (ii) both $CEAF_e$ recall and precision significantly decrease; and (iii) BLANC recall notably decreases so that F1 drops significantly in comparison to Base” (Moosavi and Strube, 2016).

As a consequence, they propose *LEA*, a new method supposed to overcome the previous limitations. *LEA* (Moosavi and Strube, 2016; Moosavi, 2020) represents the importance of each entity through a weight (meaning that *LEA* considers that a spurious singleton is less harmful than a spurious long entity). Since *LEA* is link-based, the metric is not affected by the mention-identification effect (Moosavi and Strube, 2016), which causes mention-based metrics scores to be less interpretative. Finally, even if *LEA* gives a representation to singletons (i.e. a self-link), a singleton is considered to be found only if the singleton is correctly resolved.

Despite these improvements, it has been demonstrated that *LEA* also suffers from a number of problems (Lion-Bouton et al., 2020). More specifically, *LEA* does not support specific configurations (esp. what the authors of (Lion-Bouton et al., 2020) call ‘triangular inequality’, which is not a rare case). Very similar configurations may lead to very different scores, which means systems cannot be compared and results cannot be directly interpreted.

3. The Inadequacy of Traditional Evaluation Schemes for Literary texts

In this section, we observe the application of the previous measures (that constitute de facto standard, despite their limitations) to the analysis of long texts, especially fictions and novels.

3.1. The Inadequacy of Reference Corpora

Coreference is usually evaluated on a specific corpus called OntoNotes. OntoNotes was built in the 2000s by a group of researchers from different American universities and covering a number of structural annotations (Hovy et al., 2006). This corpus has been especially popular for coreference evaluation, as it is one of the very rare resources available for the task. It gathers different kinds of texts (although they are mostly news) and is thus considered quite balanced and representative.

However, the OntoNotes corpus suffers from a number of shortcomings. The corpus is made of short texts (including short extracts of longer texts), which means coreference chains are local and limited in size. The task is thus limited in its complexity, for example the set of

antecedents for a given pronoun is small, given the short size of the text to consider. Coreference is also centered around pronouns, whereas real texts provide a wider diversity of coreferences, especially definite descriptions (“the lady”, “the milkman”).

Long texts, especially novels, are of course very different. Table 1 reports some numbers that show a clear difference between OntoNotes and a novel (*Manon Lescaut* in this case – numbers would be different with other novels but the same tendency could be observed).

	Ontonotes (average over 3,493 documents)	Manon Lescaut
# of mentions	56	11,888
# of chains	13	143
# of mentions per Chains (incl. singletons)	-	59
# of mentions per Chains (without singleton)	4.1	83
# of singletons	0	56
# of tokens	466	77,232

Table 1
Comparison of Ontonotes with a typical French novel, *Manon Lescaut*.

As one can see, the differences are huge. While OntoNotes documents are made of a few short chains (4.1 mentions per chain on average), novels like *Manon Lescaut* have 83 mentions per coreference chain on average. Moreover the distribution of chain lengths is very skewed, as nearly half of the mentions in *Manon Lescaut* correspond to one character only (the so called *Chevalier Des Grieux*), see Figure 1.




Figure 1: The distribution of coreference chains in *Manon Lescaut*, based on the number of mentions per chain. As one can see, the distribution is not even Zipfian (in which case, the curve would be straight), the number of very long chains being extremely limited. Experiments with other novels, by Balzac among others, show a quasi Zipfian curve, showing that more characters with first and secondary roles appear in Balzac’s world.

Even if *Manon Lescaut* is thus an extreme case, most novels are centered around a handful of characters that correspond to most of the mentions. Lastly, one may notice differences due to the annotation schemes, as OntoNotes does not consider singletons for the coreference task. As a conclusion, we think that even if OntoNotes is a *de facto* standard for coreference evaluation, it cannot be used as an indicator of future performance on the fiction domain.

3.2. The Low Readability of Existing Evaluation Methods

Besides the known limitations of existing evaluation methods that we have seen in section 2, these metrics are not ideal in that they all give different scores for a same input (but these differences are not directly interpretable). For example, Pradhan et al. (Pradhan et al., 2014) describe precisely how these metrics work on a toy example. In the end, their example corresponds to the following scores (F-measure): 0.4 for MUC, 0.46 for B3, 0.52 for CEAF and 0.36 for Blanc (so 0.16 difference between Blanc and CEAF). Given that different choices can be made concerning the way to calculate precision and recall (mention-based or relation-based, taking into account chain length or not, etc.) it is normal that each method has different scores for precision and recall. However, this remains a major issue.

The problem is exacerbated in the case of long novels, as one can imagine. Other issues concern the reference: some corpora/experiments consider only pronouns, some others also definite descriptions (like “the coachman”), which should be part of the task but are harder to annotate and evaluate. For example, Bamman et al. (Bamman et al., 2020) say that “OntoNotes and LitBank do not link two generic common nouns, while PreCo does” (OntoNotes, LitBank and PreCo are three corpora used in Bamman et al.’s paper), which makes the comparison between the different results delicate. The annotation of singletons also differs a lot from one corpus to the other. All these are known problems, that have reinforced the status of OntoNotes as a de facto standard (since OntoNotes provides stability, although the corpus is limited and not always relevant).

Given all these differences, a known strategy consists in calculating the mean of the different evaluation scores (or of a subset of these scores – for example (Denis and Baldrige, 2009) propose to combine MUC, CEAF and B3), a score sometimes called the CoNLL score (since it was used at CoNLL 2012 (Pradhan et al., 2012)). In our opinion, this is probably worse than keeping the different scores apart: since these scores have few in common, calculating the mean does not really make any sense (like calculating the mean of an heterogeneous set of elements).

These scores may be useful in an NLP perspective, if one really needs or wants to compare systems (although we, and many others, have shown that this comparison is dubious). But they are probably useless in a Digital Humanities perspective (or just from a linguistic point of view), where one needs to have a clear view of what is going on. In this context, the key question is to know what kinds of coreference chains are accurately analyzed, and which ones may suffer from errors.

4. Proposals for a Better Evaluation of Coreference in Long Documents

Following what we have shown in the previous sections, we think that a unique score is unable to take into account all the complexity of coreference. Moreover, a unique score (or even a combination of scores) is not really useful in the context of fictions, since no score is really interpretable in itself and scores are not comparable with each other.

It would thus make more sense to promote indicators based on the linguistic content that is measured. In the case of fictions for example, at least three different cases should be considered:

- long coreference chains, especially those including a named entity (i.e. where one of the mentions is a named entity). This generally corresponds to main characters and is especially important for most analyses (esp. representing the plot).
- singletons. We have seen that most metrics have a hard time with singletons, since by definition a singleton (a single mention) is not a chain (that requires at least two mentions linked together). Hence the choice of MUC (that ignores singletons) or LEA (that consider singleton as an entity with a link to itself). Singletons are important to evaluate the “population in a novel” (the number of characters in a novel) but should be considered in isolation and probably not as a coreference element (of course, a singleton can be included in a longer coreference chain by mistake, but in such a case, it would just be a mistake that is independent from the singleton issue).
- remaining coreference chains refer for the most part to secondary characters or just ‘background characters’, who are generally not named but can be important for the plot. Being able to evaluate the accuracy of the analysis on this kind of character is thus key. It can also be problematic, as the issue is not always fully decidable. For example, if there are different mentions of ‘coachman’ in a novel (referred to just as ‘coachman’), should we consider that they all refer to the same character, or not? This kind of precision is probably not possible in a number of cases in each novel, which means that calculating precisely the “population of a novel” is not possible with perfect accuracy (Woloch, 2004; Lavocat, 2020).

Traditional scores can then be calculated for the different cases and thus give a better picture of the situation at stake. This method may make it more difficult to compare different systems, but existing measures do not allow such a comparison anyway. A more fine-grained analysis based on the linguistic content would probably be more useful for digital humanists than current NLP-based way of reporting results.

For example, when analyzing the BookNLP¹ scores on the French LitBank², it was observed that long coreference chains were sometimes broken in different parts, thus showing the necessity to reinforce the clustering dimension of the algorithm (this procedure makes it possible to calculate coreference over different mentions of a same entity, thus reinforcing the quality of the analysis despite surface variation).

The same strategy can be used to track the analysis of singletons for example. This is in line with the approach described in (Bamman et al., 2020), where the authors propose some manipulations and an evaluation with the B^3 metric to study precisely the treatment of singletons with their approach (section 6.1, page 51 of (Bamman et al., 2020)). The same approach can be done on the French LitBank to refine this part of the analysis.

5. Conclusion

In this short paper, we presented known issues with evaluation scores related to coreference, and we have shown that this inadequacy is especially meaningful for long texts. We have

¹<https://www.lattice.cnrs.fr/projets/booknlp/>

²<https://github.com/lattice-8094/fr-litbank>

proposed a new way of evaluating coreference, taking into account the application context (digital humanities) as well as the need to have more informative and interpretable metrics. In the future, we will apply this approach to more novels, in different languages, to ensure the approach gives an appropriate representation of the problem, help the user spot weaknesses in her/his models and, above everything, ensure that the metrics used are both readable and interpretable.

References

- M. Poesio, J. Yu, S. Paun, A. Aloraini, P. Lu, J. Haber, D. Cokal, Computational models of anaphora, *Annual Review of Linguistics* 9 (2023) 561–587. URL: <https://qmro.qmul.ac.uk/xmlui/handle/123456789/83033>. doi:10.1146/annurev-linguistics-031120-111653.
- D. Bamman, T. Underwood, N. A. Smith, 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)*, Association for Computational Linguistics, Baltimore, Maryland, 2014, pp. 370–379. URL: <https://aclanthology.org/P14-1035>. doi:10.3115/v1/P14-1035.
- R. Grishman, B. Sundheim, Message understanding conference-6: A brief history, in: *Proceedings of the 16th Conference on Computational Linguistics - Volume 1, COLING '96*, Association for Computational Linguistics, Stroudsburg, PA, USA, 1996, pp. 466–471. URL: <http://dx.doi.org/10.3115/992628.992709>. doi:10.3115/992628.992709.
- M. Borovikova, L. Grobol, A. Halftermeyer, S. Billot, A methodology for the comparison of human judgments with metrics for coreference resolution, in: *Proceedings of the 2nd Workshop on Human Evaluation of NLP Systems (HumEval)*, Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 16–23. URL: <https://aclanthology.org/2022.humeval-1.2>. doi:10.18653/v1/2022.humeval-1.2.
- A. Lion-Bouton, L. Grobol, J.-Y. Antoine, S. Billot, A. Lefeuvre-Halftermeyer, Comment arpenter sans mètre : les scores de résolution de chaînes de coréférences sont-ils des métriques ? (do the standard scores of evaluation of coreference resolution constitute metrics ?), in: *Actes de la 6e conférence conjointe Journées d'Études sur la Parole (JEP, 33e édition), Traitement Automatique des Langues Naturelles (TALN, 27e édition)*, ATALA et AFPCP, Nancy, France, 2020, pp. 10–18. URL: <https://aclanthology.org/2020.jeptalnrecital-eternal>.
- N. S. Moosavi, M. Strube, Which coreference evaluation metric do you trust? a proposal for a link-based entity aware metric, in: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Berlin, Germany, 2016, pp. 632–642. URL: <https://aclanthology.org/P16-1060>. doi:10.18653/v1/P16-1060.
- N. S. Moosavi, *Robustness in Coreference Resolution*, University of Heidelberg, 2020. URL: <http://www.ub.uni-heidelberg.de/archiv/27919>. doi:10.11588/heidok.00027919.
- M. Vilain, J. Burger, J. Aberdeen, D. Connolly, L. Hirschman, A model-theoretic coreference scoring scheme, in: *Sixth Message Understanding Conference (MUC-6): Proceedings of a Conference Held in Columbia, Maryland, November 6-8, 1995*, 1995, pp. 45–52. URL: <https://aclanthology.org/M95-1005>.
- M. Recasens, E. Hovy, Blanc: Implementing the rand index for coreference evaluation, *Nat.*

- Lang. Eng. 17 (2011) 485–510. URL: <https://doi.org/10.1017/S135132491000029X>. doi:10.1017/S135132491000029X.
- X. Luo, S. Pradhan, M. Recasens, E. Hovy, An extension of BLANC to system mentions, in: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), Association for Computational Linguistics, Baltimore, Maryland, 2014, pp. 24–29. URL: <https://aclanthology.org/P14-2005>. doi:10.3115/v1/P14-2005.
- A. Bagga, B. Baldwin, Algorithms for scoring coreference chains, in: The first international conference on language resources and evaluation workshop on linguistics coreference, volume 1, 1998, pp. 563–566.
- X. Luo, On coreference resolution performance metrics, in: Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Vancouver, British Columbia, Canada, 2005, pp. 25–32. URL: <https://aclanthology.org/H05-1004>.
- E. Hovy, M. Marcus, M. Palmer, L. Ramshaw, R. Weischedel, OntoNotes: The 90% solution, in: Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers, Association for Computational Linguistics, New York City, USA, 2006, pp. 57–60. URL: <https://aclanthology.org/N06-2015>.
- S. Pradhan, X. Luo, M. Recasens, E. Hovy, V. Ng, M. Strube, Scoring coreference partitions of predicted mentions: A reference implementation, in: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), Association for Computational Linguistics, Baltimore, Maryland, 2014, pp. 30–35. URL: <https://aclanthology.org/P14-2006>. doi:10.3115/v1/P14-2006.
- D. Bamman, O. Lewke, A. Mansoor, An annotated dataset of coreference in English literature, in: Proceedings of the Twelfth Language Resources and Evaluation Conference, European Language Resources Association, Marseille, France, 2020, pp. 44–54. URL: <https://aclanthology.org/2020.lrec-1.6>.
- P. Denis, J. Baldridge, Global joint models for coreference resolution and named entity classification, *Procesamiento del lenguaje natural* 42 (2009).
- S. Pradhan, A. Moschitti, N. Xue, O. Uryupina, Y. Zhang, CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes, in: Joint Conference on EMNLP and CoNLL - Shared Task, Association for Computational Linguistics, Jeju Island, Korea, 2012, pp. 1–40. URL: <https://aclanthology.org/W12-4501>.
- A. Woloch, *The one and the Many, Minor characters and the Space in the Novel*, Princeton University Press, Princeton, 2004.
- F. Lavocat, *Les personnages rêvent aussi*, Hermann, coll. "Fictions pensantes", Paris, 2020.