

Towards Multiple Antecedent Coreference Resolution in Specialized Discourse

Alicia Burga¹, Sergio Cajal¹, Joan Codina-Filbà¹, Leo Wanner^{1,2}

¹NLP Group, Department of Information and Communication Technologies
Pompeu Fabra University
Roc Boronat, 138, 08018 Barcelona, Spain
alicia.burga, sergio.cajal, joan.codina, leo.wanner@upf.edu

²Catalan Institute for Research and Advanced Studies (ICREA)

Abstract

Despite the popularity of coreference resolution as a research topic, the overwhelming majority of the work in this area focused so far on single antecedent coreference only. Multiple antecedent coreference (MAC) has been largely neglected. This can be explained by the scarcity of the phenomenon of MAC in generic discourse. However, in specialized discourse such as patents, MAC is very dominant. It seems thus unavoidable to address the problem of MAC resolution in the context of tasks related to automatic patent material processing, among them abstractive summarization, deep parsing of patents, construction of concept maps of the inventions, etc. We present the first version of an operational rule-based MAC resolution strategy for patent material that covers the three major types of MAC: (i) nominal MAC, (ii) MAC with personal / relative pronouns, and MAC with reflexive / reciprocal pronouns. The evaluation shows that our strategy performs well in terms of precision and recall.

Keywords: Multiple antecedent coreference, split antecedent coreference, coreference evaluation, patent processing

1. Introduction

Coreference resolution has been a popular research topic for a considerable time now both in theoretical and in computational studies; see, among others, (Lasnik, 1989; L.Eguren and Fernández Soriano, 2007; Kayne, 2005) for the first and (Mitkov, 1999; Recasens et al., 2010; Recasens and Hovy, 2010) for the latter. It is an obligatory task for any language understanding application that goes beyond surface-oriented parsing. However, the overwhelming majority of the state-of-the-art coreference resolution works focused so far exclusively on one kind of coreference, namely the *single* antecedent coreference illustrated in (1):

- (1) [John]_i met Mary in New York. At that time, [he]_i studied Computer Science.

Multiple antecedent coreference (henceforth MAC), or *split antecedent coreference* (Lasnik, 1989; Hornstein, 1999), as illustrated in the examples (2) and (3), has been largely neglected.

- (2) [John]_i convinced [Mary]_j to move to Barcelona. Now, [both]_{i+j} enjoy living close to the sea.
- (3) [John]_i, who adores [Mary]_j from the first day, thinks that [they]_{i+j} make a great couple.

Some computational proposals that deal with coreference mention the phenomenon, but do not treat it (and do not justify in a satisfactory way the decision to ignore it either); see, e.g., (Martínez-Barco and Palomar, 2011; Baldwin, 1997). Others treat it, but do not present the achieved performance in detail (Aone and Bennett, 1995; Kennedy and Boguraev, 1996). But, as already pointed out, most of the existing proposals (among them, e.g., (Ge et al., 1998;

Refoufi, 2007)) do not even mention the phenomenon. The exclusion of MAC from coreference resolution is also reflected in the annotation tools that are frequently used: either they are not implemented to directly capture this phenomenon, as, e.g., GATE (Cunningham et al., 2011), or Clinka (Orasan, 2000), or they are partially implemented to handle it, but presumably require an advanced user configuration for this purpose, as, e.g., MMAX (Müller and Strube, 2006), or their specification does not mention the phenomenon at all, as, e.g., in the case of AncoraPipe (Bertrán et al., 2008).

The most obvious explanation for this gap is the implicit assumption that MAC is scarce in general discourse. And indeed, a quick examination of two small newspaper articles (in total 1354 tokens)¹ seems to buttress this assumption. Even if plural coreferential elements appear in these articles with a certain frequency (31 times), we do not find any multiple antecedence: the identified plural coreferential elements are NPs with one single (plural) antecedent, as in the following example:

- (4) What about the sacrifices of [those people]_i, did [they]_i forget about that?

Cursory checks make us hypothesize that MAC is more common in spoken general discourse, but a more thorough study would be required to make any firm statement in this respect. In any case, the situation is drastically different in specialized discourse such as patents. In patent material, MAC is much more frequent than in general domain.

¹<http://www.express.co.uk/news/world/609410/refugee-crisis-Germany-fights-religion-Kassel-Rainer-Wendt-police-migrants>; <http://www.usnews.com/news/articles/2014/08/26/islamic-state-calls-for-revenge-against-syrias-rebels>

Consider, for illustration, a typical example copied from a patent:

- (5) *The electric circuit wherein each of the DC-to-AC converters comprises [a first switch]_i [...] and [a second switch]_j [...] The electric circuit wherein [the first and second switches]_{i+j} are connected [...]*

This frequency is high to an extent that makes the automatic resolution of MAC in patents a precondition for successful automatic identification of the components of an invention (and thus for the content-oriented analysis of patents—for instance, in the context of a more global task of the construction of a complete conceptual representation of an invention). Thus, to correctly represent the content of the statement *The control circuit controls the two batteries* in semantic terms of the kind `control(control_circuit, first_battery) ∧ control(control_circuit, second_battery)`, we need to resolve the MAC between *first battery*, *second battery* and *batteries*.

In what follows, we present a proposal for the resolution of the prominent types of MAC identified in patent material. In the next section, we introduce the phenomenon of MAC in more detail. Section 3 outlines our proposal for MAC resolution. Section 4 discusses the problem of the evaluation of the performance of MAC resolution models and presents the outcome of the evaluation of our model. Section 5, finally, draws some conclusions from the presented work and outlines the future research that we plan to carry out in order to further advance in the area of MAC resolution in patent material.

2. Types of multiple antecedence

In an empirical study of patent material, we identified three major types of multiple antecedent coreference, which also coincide with the three types of referential expressions discussed in the context of the Government and Binding Theory (Chomsky, 1993; Haegeman, 1994):

Nominal: In this type of multiple antecedence, all involved elements have nominal heads. In most cases, all elements share the same head, although the non-antecedent element always carries plural; cf. (6). However, each NP may also have a different head; cf. (7). The multiple antecedents can be singular or plural.

- (6) *[The math teacher]_i had lunch with [the history teacher]_j. [Both teachers]_{i+j} are old friends.*

- (7) *[Lisa]_i smiled at [Bart]_j. [The children]_{i+j} were happy.*

Given its legal and descriptive nature, the language of patents seeks to be as explicit as possible. Therefore, in the Nominal MAC, which is the most common multiple antecedent coreference in patents, the head is shared by all elements, but the antecedents are modified by their own ordinal and/or contrastive adjectives as in (5) above.²

²We refer to those adjectives as “contrastive adjectives” that are complementary to each other; consider, e.g., *top* and *bottom*, *positive* and *negative*, *anterior*, *middle* and *posterior*, etc.

In general, we can state that a plural non-pronominal NP is a candidate to have more than one antecedent if the potential antecedents are outside of its sentence.³ Therefore, to cover this subcase of MAC, inter-sentential search for antecedents must be implemented.

Pronominal: In this type of multiple antecedence, the element that corefers with more than one antecedent is pronominal. The pronoun can be personal or relative, as in (9) and (10), correspondingly.⁴

- (9) *[Lisa]_i smiled at [Bart]_j. [They]_{i+j} were happy.*

- (10) *[Lisa]_i and [Bart]_j, [who]_{i+j} are siblings, like watching TV together.*

A pronoun is thus a candidate to belong to a multiple antecedent configuration if there are at least two (pro)nominal phrases outside of its clause (i.e., the antecedents on one hand, and the pronoun on the other hand, are headed by different verbs). In the case of personal pronouns, the potential antecedents can be in the same sentence or in a different one, but in the case of relative clauses, they are always in the same sentence.⁵

Reciprocal/reflexive: This subtype of multiple antecedence involves reflexive and reciprocal pronouns (e.g., *themselves* and *each other*, respectively). When these pronouns have more than one antecedent, those are NPs. Different from previous cases, though, the antecedents always belong to the same clause as the reciprocal/reflexive pronoun.⁶

- (11) *[Lisa]_i and [Bart]_j smiled at [each other]_{i+j}.*

A reflexive/reciprocal pronoun is a candidate to have multiple antecedents if there are at least two NPs inside its clause (i.e., the NPs are governed by the same verb as the reciprocal/reflexive pronoun).⁷

3. Implementation of MAC resolution

Single coreference resolution strategies can be roughly grouped into those that treat coreference as a set of binary links between mentions (such that coreference holds between two subsequent individual mentions of a concept),

³This requirement is directly related to Principle C of Government and Binding, which establishes that a referential expression must be free everywhere. Given that it is possible to find a coreference between NPs of two sentences coordinated by a conjunction, we restrict our definition of sentence to the unit with just one main verb.

⁴Even if cases that do not involve coordination, such as, e.g., in *The [charger]_i is assembled on top of the [battery]_j. [They]_{i+j} ...*, should also be considered in the context of Pronominal MAC, they are neither described nor taken into account in our study because they do not appear in patents.

⁵This requirement correlates with the Principle B of Government and Binding: a pronoun must be unbound in its domain.

⁶In the patent domain, the presence of reflexive pronouns is very limited. On the other hand, there are many reciprocal pronouns.

⁷Cf. Principle A of Government and Binding: an anaphora (reflexive/reciprocal pronouns) must be bound in its domain.

and those that treat it as a cluster of entities (such that coreference holds between all mentions of a concept). In both, all mentions of the same concept across the text are captured: in the first, in terms of a concatenation of binary links between mentions into a chain (or, more generally, into a connected graph); in the second, in terms of a bag of mentions. In contrast, MAC is a local phenomenon: the relation between a given referent and its antecedents is unique in the sense that a MAC antecedent cannot act in its turn as referent—even if the referent or one or more of the antecedents obviously can also belong to distinct single coreference chains. When designing the multiple antecedence coreference resolution strategy, this locality should be taken into account.

Below, we outline our multiple antecedence coreference resolution strategy in patents for all three major types of MAC (see the previous section). The NPs that serve as antecedents in all three types can be assumed to contain a head, modified by an ordinal or a contrastive adjective ($Adj_{s,2}$) and an optional non-contrastive adjective (Adj_a); if both modifiers appear, the contrastive adjective is more external, and therefore it appears further from the head:

$$(12) NP_1: Adj_1 + (Adj_a) + Head_{sg}^i \text{ }^8$$

$$(13) NP_2: Adj_2 + (Adj_a) + Head_{sg}^i$$

The strategy has thus to relate antecedents of this kind with their referent, which is specific to each type of MAC.

The strategy has been implemented in Java as a GATE (Cunningham et al., 2011) processing resource plug-in. The plug-in receives as input nominal chunks of the syntactic dependency tree of a patent sentence, obtained using the joint morphological and syntactic parser (Bohnet et al., 2013) from MATE tools.⁹ and outputs the list of MACs. As in Stanford Coreference Resolution (Raghunathan et al., 2010), we solve each case of multiple antecedence coreference as an independent rule-based sieve.¹⁰

3.1. Nominal multiple coreference

The resolution of multiple antecedent coreference between nominal elements is carried out in two stages: 1. *Antecedent grouping*, and 2. *Antecedents and plural match*. During the antecedent grouping stage, singular NPs that appear in the same sentence and share number, head and pre-nominal modifiers (except those that are ordinal or contrastive) are grouped as a single element. Thus, for instance, if there are two NPs that follow the pattern detailed in (12) and (13), such as *first alkaline battery* and *second alkaline battery*, the system groups both NPs into a single plural NP_3 , getting the phrase *first and second alkaline batteries*, which has the following structure:

$$(14) NP_3: Adj_1 + Adj_2 + (Adj_a) + Head_{pl}^i$$

⁸The parentheses mark optional elements.

⁹The parser has been optimized for parsing patent material (Pekar et al., 2014).

¹⁰For single antecedent coreference resolution in patents that uses an adapted version of Stanford Coreference Resolution, see (Bouayad-Agha et al., 2014).

After the grouping stage, the second stage carries out the matching. During this second stage, each constructed NP of the kind of NP_3 , resulting from the first stage, is matched with the closest posterior (with respect to NP_1 and NP_2) plural NP that has the same head and pre-nominal modifiers as NP_1 and NP_2 . When establishing the match, the following criteria are used:

- (i) Non-contrastive pre-nominal modifiers are optional in the referential phrase; if such modifiers are present, they must match the modifiers in the antecedent NPs (e.g., *the top battery charger* [...] *the bottom battery charger* can corefer either with *the chargers*, *the battery chargers*, or with *the top and bottom (battery) chargers*).
- (ii) Contrastive adjectives are also optional in the referential phrase; if they are present, all of them must equally appear in the antecedent NPs. Thus, there is a multiple antecedent coreference between *the first battery* [...] *the second battery* [...] and *the first and second batteries*, but not between *the first battery* [...] *the second battery* [...] and *the first batteries*.

So, with the plural NP at hand, the system looks for possible matches (according to the criteria described), as in simple coreference. It establishes a match when detected, and then it splits the antecedent into its multiple components, to strictly solve MAC. Thus, we can describe the process of resolving nominal MAC in terms of the following four sequential steps:

1. **Composition:** $(NP_1 + NP_2) = NP_3$
2. **Matching:** $NP_3 \Leftrightarrow NP_x$
3. **Decomposition:** $NP_3 = (NP_1 + NP_2)$
4. **MAC resolution:** $(NP_1 / NP_2) \Leftrightarrow NP_x$

Most cases of nominal MAC are anaphoric: the antecedent NPs precede the referential element. However, there are also cases of cataphora, where the coreferential element precedes the antecedent(s). In order to cover both anaphora and cataphora, the sieve is executed in two runs: from the beginning to the end of the text, and from the end to the beginning of the text.

3.2. Multiple coreference involving personal/relative pronouns

Multiple antecedent anaphora resolution (i.e., coreference resolution that involves pronouns) varies depending on the type of pronoun involved.

The anaphora resolution sieve that covers plural personal pronouns consists of two steps. First, the pronouns are identified.¹¹ Second, from the pronoun location, the system goes backwards over the text until it finds two or more coordinated NPs (each having the structure of NP_1 and NP_2 in (12) and (13) above) and assigns them the status of antecedents of the pronoun; cf. the following example:

¹¹In patents, the only plural pronouns that appear are those in third person: *they* and *them*.

(15) [...] comprising [a top battery]_{A₁} and [a bottom battery]_{A₂}. They are connected [...]

Since our search for antecedents is restricted to coordinated NPs, cases where there is no explicit coordination (through a conjunction) are not detected (as, e.g., *The first battery is connected to the starting device. The second battery is in reserve. Both are alkaline.*).

The MAC resolution sieve that includes relative pronouns is carried out in the same way as the one involving personal pronouns, but the search is limited to the same sentence as the relative pronoun; cf. (16):

(16) [...] comprising [a top battery]_{A₁} and [a bottom battery]_{A₂}, which are connected [...]

3.3. Reflexive/reciprocal pronoun multiple reference

The nature of multiple coreference anaphora resolution that involves reflexive/reciprocal pronouns implies the use of syntactic parsing for finding the multiple antecedents, given that antecedents and pronoun should depend on the same head. However, as the parser performance is not optimal in patents (e.g., it does not always detect correctly coordinations or the scope of the clause) (Burga et al., 2013), instead of relying on the resulting dependency tree, we use other criteria to establish the clause and, therefore, restrict the search: the system only evaluates cases in which there is at most one content verb between the pronoun and the antecedent candidate (with the candidates having the same structure as NP₁ and NP₂ in (12) and (13) above), and where there are no determined punctuation marks (“:” and “;”) between them. Also, the system considers as antecedent candidates only those NPs that are not prepositional complements. An example of this subtype of MAC is shown below in (17):

(17) [...] [a top battery]_{A₁} and [a bottom battery]_{A₂} are connected to each other [...]

4. Evaluation of MAC Resolution

4.1. General Considerations

The evaluation metrics used in single coreference resolution strategies reflect the two different views on coreference (as a connected graph of mentions or as a bag of mentions) mentioned in Section 3. For the first, see, e.g., MUC and Pairwise F1 (Vilain et al., 1995) or Rand (Finkel and Manning, 2008); for the latter, e.g., B-cube (Bagga and Baldwin, 1998), CEAF (Luo, 2005), ACE-Value (NIST, 2003), Mutual Information (Popescu-Belis et al., 2004), BLANC (Recasens and Hovy, 2011). As already in the case of coreference resolution, we cannot use the same metrics for MAC resolution.

In accordance with the locality of MAC, the evaluation of MAC resolution strategies can be interpreted as a two step task:

- (i) assessment of the quality of a strategy with respect to the identification of the referent: how many of the MAC referents in the ground truth have been detected and how many mentions have been erroneously considered to be MAC referents; and

- (ii) assessment of the quality of a strategy with respect to the antecedents of a given referent: how many of the antecedents of each referent have been recognized and, again, how many mentions have been considered erroneously to be antecedent of a MAC referent.

For both, we use precision and recall. The assessment of the quality of antecedent recognition is done on correctly recognized referents only.

Obviously, we could have measured only precision and recall of the detection of antecedents, without considering whether the referent has been recognized correctly or not, and possibly even combining precision and recall into a single *F₁ score*. However, while this would give us a suitable measure to compare the performance of different approaches, it would not provide us with any details with respect to the behavior of our approach.

4.2. Evaluation of MAC resolution on patent material

To evaluate the proposed MAC resolution model on patent material, we manually annotated a patent corpus with MAC. To ensure a sufficiently large ground truth, the annotated corpus consists of sentences that not necessarily form a coherent chunk of text.

For evaluation, the corpus was first parsed with Bohnet et al. (2013)’s parser, which was further adapted to the patent domain (Pekar et al., 2014). To assess the influence of parser errors on the quality of our MAC resolution model, we ran two experiments. In the first, we used the output of the parser as is; in the second, we manually corrected the parser output and then ran the MAC module on the corrected output.

The results are shown in Table 1. ‘referent precision’ / ‘referent recall’ reflects how the quality of the detection of the referents in the gold dataset; ‘antecedent precision’ / ‘antecedent recall’ assesses the quality of the identification of antecedents attached to a correctly identified referent only; ‘global precision’ and ‘global recall’ as well as the corresponding *F₁ score* reflect the performance of our model on all antecedents, no matter whether their referents have been recognized correctly or not.

It can be observed that we achieve a global *F₁ score* of 80%, with precision higher than recall. With respect to the parser error propagation, we see a drop of 7.5% in the final *F₁ score*, compared to the use of corrected parser output. In this case, most of the drop affects the recall (10%), while precision decreases by a 1.9% only. In other words, referent detection is more sensible to parser errors than the detection of antecedents.

5. Conclusions

The phenomenon of multiple antecedence coreference is prominent in patent material, such that it appears necessary to develop a MAC resolution model in order to provide the necessary information for such downstream tasks as information extraction from patents, abstractive summarization of patents, invention-oriented conceptual map construction, etc.

run	referent precision	referent recall	antecedent precision	antecedent recall	global precision	global recall	global F_1
uncorrected parse	89.4%	77.3%	90.6%	85.3%	85.2%	63.0%	72.5 %
corrected parse	95.0%	86.3%	91.8%	87.1%	87.1%	73.9%	80.0 %

Table 1: Evaluation results of the proposed MAC resolution model on the original and corrected parser output

We presented a rule-based sieve-oriented proposal to handle the three major types of multiple antecedence coreference resolution. Our model achieved good performance figures (85.2% precision and 63.0% recall on uncorrected parser output, without any restriction to correctly recognized referents), although the recall can (and should) be improved. The relatively low (compared to precision) recall can be explained by the paradigm we adopted for our model: it is rule-based and, therefore, its coverage is limited to the scope of the rules. Cases that are not covered by the rules are not recognized. On the other hand, the annotation of patent material with coreference phenomena is a very time consuming and complex task, such that the compilation of a sufficiently large training dataset for a statistical approach appears out of the reach—at least for a timely constrained exercise. The expansion of the rule set for higher coverage seems thus to be the most obvious solution. In the future, we will dedicate further work to this task.

Our experiments have also shown that erroneous parser output significantly affects the quality of MAC resolution. Since MAC resolution is bound to use syntactic features (be it for training or in rule conditions), additional work must also be invested into the amelioration of the parser performance on patent material.

6. Acknowledgements

The research reported in this paper has been carried out in the scope of the iPatDoc project, which was funded by the European Commission under the contract number FP7-SME-606163.

7. Bibliographical References

- Aone, C. and Bennett, S. (1995). Evaluating Automated and Manual Acquisition of Anaphora Resolution Strategies. In *Proceedings of the 33rd Annual Meeting on Association for Computational Linguistics, ACL '95*, pages 122–129, Stroudsburg, PA, USA.
- Bagga, A. and Baldwin, B. (1998). Algorithms for scoring coreference chains. In *The first international conference on language resources and evaluation workshop on linguistics coreference*, volume 1, pages 563–566. Citeseer.
- Baldwin, B. (1997). Cogniac: high precision coreference with limited knowledge and linguistic resources. In *Operational Factors in Practical, Robust Anaphora Resolution for Unrestricted Texts*, pages 38–45.
- Bertrán, M., Borrega, O., Recasens, M., and Soriano, B. (2008). AnCoraPipe: A tool for multilevel annotation. *Procesamiento del Lenguaje Natural*, 41.
- Bohnet, B., Nivre, J., Boguslavsky, I., Farkas, R., Ginter, F., and Hajić, J. (2013). Joint morphological and syntactic analysis for richly inflected languages. *Transactions of the Association for Computational Linguistics*, 1:415–428.
- Bouayad-Agha, N., Burga, A., Casamayor, G., Codina, J., Nazar, R., and Wanner, L. (2014). An exercise in reuse of resources: Adapting general discourse coreference resolution for detecting lexical chains in patent documentation. In *Proceedings of the International Conference on Linguistic Resources and Evaluation (LREC)*, Reykjavik, Iceland.
- Burga, A., Codina, J., Ferraro, G., Saggion, H., and Wanner, L. (2013). The challenge of syntactic dependency parsing adaptation for the patent domain. In *Proceedings of ESSLLI*.
- Chomsky, N. (1993). *Lectures on Government and Binding: The Pisa lectures*, volume 9. Walter de Gruyter.
- Cunningham, H., Maynard, D., Bontcheva, K., Tablan, V., Aswani, N., Roberts, I., Gorrell, G., Funk, A., Roberts, A., Damljanovic, D., Heitz, T., Greenwood, M. A., Saggion, H., Petrak, J., Li, Y., and Peters, W. (2011). *Text Processing with GATE (Version 6)*.
- L. Eguren et al., editors. (2007). *Coreference, Modality and Focus*. John Benjamins.
- Finkel, J. R. and Manning, C. D. (2008). Enforcing transitivity in coreference resolution. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies: Short Papers*, pages 45–48. Association for Computational Linguistics.
- Ge, N., Hale, J., and Charniak, E. (1998). A statistical approach to anaphora resolution. In *Proceedings of the Sixth Workshop on Very Large Corpora*, pages 161–170.
- Haegeman, L. (1994). *Introduction to Government and Binding Theory*. Blackwell Textbooks in Linguistics. Wiley.
- Hornstein, N. (1999). Movement and control. In *Linguistic Inquiry*, pages 30–69.
- Kayne, R. (2005). *Movement and Silence*. Oxford University Press.
- Kennedy, C. and Boguraev, B. (1996). Anaphora for Everyone: Pronominal Anaphora Resolution Without a Parser. In *Proceedings of the 16th Conference on Computational Linguistics - Volume 1*, pages 113–118.
- Lasnik, H. (1989). *Essays on Anaphora*. Kluwer Academic Publishers.
- Luo, X. (2005). On coreference resolution performance

- metrics. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, pages 25–32. Association for Computational Linguistics.
- Martínez-Barco, P. and Palomar, M. (2011). Computational Approach to Anaphora Resolution in Spanish Dialogues. *CoRR*, abs/1106.0673.
- Mitkov, R. (1999). Anaphora Resolution: The State Of The Art. Technical report.
- Müller, C. and Strube, M. (2006). Multi-level annotation of linguistic data with MMAX2. In S. Braun, et al., editors, *Corpus Technology and Language Pedagogy: New Resources, New Tools, New Methods*, pages 197–214. Peter Lang.
- NIST, U. (2003). The ace 2003 evaluation plan. *US National Institute for Standards and Technology (NIST)*, pages 2003–08.
- Orasan, C. (2000). ClinkA a Coreferential Links Annotator.
- Pekar, V., Yu, J., El-karef, M., and Bohnet, B. (2014). Exploring options for fast domain adaptation of dependency parsers. In *First Joint Workshop on Statistical Parsing of Morphologically Rich Languages and Syntactic Analysis of Non-Canonical Languages*, pages 54–65, Dublin.
- Popescu-Belis, A., Rigouste, L., Salmon-Alt, S., and Romary, L. (2004). Online evaluation of coreference resolution. In *4th International Conference on Language Resources and Evaluation-LREC'04*, pages 4–p.
- Raghunathan, K., Lee, H., Rangarajan, S., Chambers, N., Surdeanu, M., Jurafsky, D., and Manning, C. D. (2010). A Multi-Pass Sieve for Coreference Resolution. In *Proceedings of the 2010 EMNLP 2010*, pages 492–501.
- Recasens, M. and Hovy, E. H. (2010). Coreference Resolution across Corpora: Languages, Coding Schemes, and Preprocessing Information. In *Proceedings of the 48th Annual Meeting of the ACL, July 11-16, 2010, Uppsala, Sweden*, pages 1423–1432.
- Recasens, M. and Hovy, E. (2011). Blanc: Implementing the rand index for coreference evaluation. *Natural Language Engineering*, 17(04):485–510.
- Recasens, M., Hovy, E., and Martí, M. A. (2010). A Typology of Near-Identity Relations for Coreference (NIDENT). In N. Calzolari (Conference Chair), et al., editors, *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta, May. European Language Resources Association (ELRA).
- Refoufi, A. (2007). A Modular Architecture for Anaphora Resolution. *Journal of Computer Science*, 3(4):199–203.
- Vilain, M., Burger, J., Aberdeen, J., Connolly, D., and Hirschman, L. (1995). A model-theoretic coreference scoring scheme. In *Proceedings of the 6th conference on Message understanding*, pages 45–52. Association for Computational Linguistics.