Multi-Sieve Pass for Coreference Resolution

Summarized by Rodrigo Agerri
IXA NLP Group
UPV/EHU
rodrigo.agerri@ehu.eus

Abstract

We introduce the multi-sieve pass for coreference resolution system developed by the NLP Group of Stanford University (Raghunathan et al., 2010), (Lee et al., 2011) and (Lee et al., 2013). This system proposes a number of deterministic passes, ranging from high precision to higher recall, each dealing with a different manner in which coreference manifests itself in running text.

1 Introduction

Sieves	Type	CONLL 2011 F1
Mention Detection	NPs, NER and PRP	-
Sieve 1	**Speaker Identification**	**29.2**
Sieve 2	**String Match**	**45.3**
Sieve 3	Relaxed String Match	45.4
Sieve 4	Precise Constructs	45.7
Sieve 5	**Strict Head Match A**	**48.5**
Sieve 6	Strict Head Match B	48.8
Sieve 7	Strict Head Match C	49.3
Sieve 8	Proper Head Noun Match	49.5
Sieve 9	Relaxed Head Match	49.7
Sieve 10	**Pronoun Match**	**59.3**

Table 1: Multisieve Pass and CoNLL 2001 F1 Evaluation

2 Mention Processing

The Mention processing module consists of the mention detection and the mention resolution architecture:

- *Mention Detection*: Noun Phrases (NP), Named Entity Recognition (NERC), Pronouns (PRP) and heuristics.
- *Resolution Architecture*:
  1. **Reduce search space**: consider only first mentions of each cluster to be resolved in the application of each sieve.
  2. **Selection of antecedents for each sieve**: left to right breadth first search.
  3. **Feature sharing**: gender, number and animacy, union for these attributes for every mention. If contradiction keep variety of attributes so that later the cluster can be linked with any mention which has any of these attributes.

2.1 Mention Detection

Their detection algorithm is based on syntactic features, NER and manually-written patterns. The manually written patterns are highly corpus-dependent, so even though we will describe them because they are used in the English coreference module, it should be noted that these patterns will probably change according to the annotation policy of each corpus. For example, the adjectival forms of nations are valid in the ACE corpus (LDC et al., 2005).

1. **Start by marking all noun phrases, pronouns and named entity mentions** (no previously marked as modifiers in the NPs).
2. **Apply the following heuristics or patterns:**
   (a) Remove a mention if a larger mention with the same head word exists. For example, remove *The five insurance companies* in *The five insurance companies approved to be established this time.*
   (b) Discard numeric entities such as percents, money, cardinals, and quantities: 9%, $10, *Tens of thousands*, *100 miles.*
   (c) Remove mentions with partitive or quantification: *a total of 177 projects*, *none of them*, *millions of people.* As a general rule, these are the NPs with 'of' preceded by one of 9 quantifiers and 34 partitives in English.
   (d) Remove pleonastic *it* pronouns, detecting using patterns such as *It is possible that ...*, *It seems that ...*, *It turns out ...*.
   (e) Discard adjectival forms of nations or nationality acronyms: *American*, *U.S.*, following Ontonotes 4.0 annotation guidelines (Weischedel et al., 2010).

(f) Remove these words: *there, ltd.,etc,'s, and hmm.*

Again, it should be stressed that **patterns are corpus-dependent**. Therefore, if we are to design these type of patterns, it should be considered that:

1. **For the first year of OpeNER**, they will need to be created for each different evaluation corpus for each language of OpeNER.

2. **For the second year**, these patterns will need to be created for each language for the evaluation datasets of reviews that need to be created.

2.2 Resolution Architecture

The resolution architecture incrementally creates mention clusters in each sieve pass application. As mentioned above, the resolution architecture at each sieve pass **(i) it filters the search space for which mention should be considered for resolution** (see section 2.2.1), **(ii) it sorts the antecedents that should be considered for resolution of a given mention** (see section 2.2.2), and **(iii) it constructs features from partially-built mention clusters** (section 2.2.3).

2.2.1 Filtering the Search Space: Mention Selection in a Given Sieve

For each sieve pass, we have partial mention clusters produced by the previous sieve pass. Of course, the first sieve pass (Exact Match) simply starts with singleton mention clusters, one for each of the mentions detected in the mention detection module (section 2.1). This information is used by:

1. Considering **only mentions** that are **first in textual order** in their cluster. Thus, if we are given the following ordered list of mentions: $m_1^1, m_2^2, m_3^2, m_4^3, m_5^1, m_6^2$, the coreference system will only try to resolve mention m_2^2 and mention m_4^3. These are the first mentions in textual order in the mention clusters to which they belong[1]. It is easy to check that mention m_1^1 does not need to be resolved because it is the first in the text, and that the rest of the mentions are not the first in mention clusters 2 and 1.

2. **Discourse Salience**: Mentions appearing *first* in their mention clusters *are not to be considered for resolution* if:

   (a) The mention starts with **indefinite pronouns** such as *some, other*, etc.

   (b) The mention starts with **indefinite articles** such as *a, an*, etc.

   (c) The mentions are **bare plurals**, such as *bars, restaurants*, etc.

EXCEPTION!!: *Discourse Salience* does **NOT APPLY** for the *Exact String Match* Sieve (section 3.2).

[1] The *superscript* indicates *mention cluster id* and the *subscript* indicates the mention id

2.2.2 Antecedent Selection for a Given Mention

Once we are given a given mention m_i, each sieve pass either will not be able to propose a solution depending on its features, or *will deterministically select a single best antecedent mention* from a list of previously ordered mentions $m_1, \ldots, m_i$. In Stanford's multi sieve pass system, the *candidate antecedents are ordered* using syntactic information from a *constituent parsing analysis*:

1. In given sentential clause, namely, those parser **constituents** which start with an **S label**, are sorted using a left-to-right breadth-first traversal of the corresponding syntactic constituent. Figure 1 shows an example of candidate ordering based on this method (Hobbs, 1978).

2. If the sentence containing the mention to be resolved contains multiple clauses, the above method is repeated separately for each $S*$ constituent, starting with the constituent in which the mention is located.

3. **Clauses in previous sentences** are sorted based on their textual proximity to the mention to be resolved.

The antecedent sorting is important because the coreference resolution system stops at the first match, namely, at the antecedent which has been placed as first by this selection method.

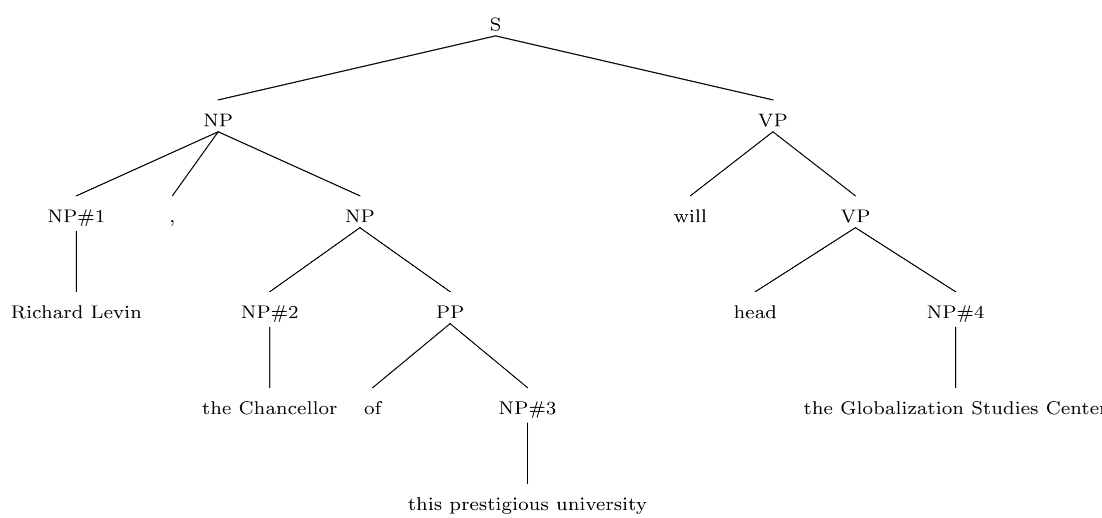

Figure 1: Left-to-right breadth first tree traversal.

2.2.3 Feature Sharing

In the multi sieve pass system, each sieve pass gets (possibly incomplete) entity[2] information for each mention located in the mention clusters built by previous sieve passes. Mentions that has not been placed in any mention cluster or entity by any of the sieve passes are singleton members of their own mention cluster. The feature sharing procedure is as follows:

1. **A union of every attribute for all mentions in a given mention cluster or entity is performed**. These attributes can be mention's number, gender, animacy, etc.

2. If attributes from two or more mentions in the same cluster contradict each other, then all attribute variants are maintained for that mention cluster or entity. For example, if two mentions in a given mention cluster display singular plural number respectively, then both attributes (singular/plural) are maintained as attributes of the mention cluster or entity.

3. Later, the mention cluster or entity displaying (singular/plural) attributes can be merged or accept further mentions that display both types of attributes.

3 Sieve Passes

3.1 Speaker Identification

This sieve matches speakers to compatible pronouns, using shallow discourse understanding to handle quotations and conversation transcripts, following the early work of Baldwin (1995, 1997). They begin by identifying speakers within text. In non-conversational text, they use a simple heuristic that searches for the subjects of reporting verbs (e.g., say) in the same sentence or neighboring sentences to a quotation. In conversational text, speaker information is provided in the dataset.

The extracted speakers then allow to implement the following sieve heuristics:

1. (I)s[3] assigned to the same speaker are coreferent.

2. (you)s with the same speaker are coreferent.

3. The speaker and (I)s are coreferent.

Thus, for example *I*, *my*, and *she* in the following sentence are coreferent: "[I] voted for [Nader] because [he] was most aligned with [my] values," [she] said. In addition, they impose speaker constraints on decisions made by subsequent sieves:

[2]Entity here refers to a cluster of mentions.

[3]They define 'I', 'my', 'me', or 'mine', (we) as first person plural pronouns, and (you) as second person pronouns.

- The speaker and a mention which is not (I) in the speakers utterance cannot be coreferent.
- Two (I)s (or two (you)s, or two (we)s) assigned to different speakers cannot be coreferent.
- Two different person pronouns by the same speaker cannot be coreferent.
- Nominal mentions cannot be coreferent with (I), (you), or (we) in the same turn or quotation.
- In conversations, you can corefer only with the previous speaker.

The constraints result in causing [my] and [he] to not be coreferent in the above example (due to the third constraint).

3.2 Exact Match

This model links two mentions only if they contain exactly the same extent text, including modifiers and determiners. For example, [the Shahab 3 ground-ground missile] and [the Shahab 3 ground-ground missile].

3.3 Precise Constructs

This pass links to mentions if any of the following conditions are satisfied:

1. **Appositive**: Two nominal mentions are in appositive construction, namely, "[Pierre Vinken], [chair of Elsevier], said...". The detection of apositives looks for the third children of a parent NP whose expansion begins with (NP, NP), when there is not a conjunction in the expansion (Haghighi and Klein, 2009).
2. **Predicate Nominative**: Two mentions, nominal or pronominal, are in copulative subject-object relation, for example, "[Pierre Vinken] is [the chairmain of Elsevier]" (Poon and Domingos, 2008).
3. **Role appositive**: the candidate antecedent is headed by a noun and appears as a modifier in a NP whose head is the current mention, for example, "[[actress] Rebecca Shaeffer]", inspired by Haghighi and Klein (2009), but further constrained as follows: This match will be considered if and only if:
   (a) The mention is labelled as a person.
   (b) The antecedent is animate (see 5.1).
   (c) The antecedent gender is not neutral.
4. **Relative Pronoun**: The mention is a relative pronoun that modifies the head of the antecedent NP: "[the finance street [which] has already formed in the Waitan district]".

5. **Acronym**: Both mentions are tagged as NNP and one of them is the acronym of the other: "[Agence France Presse], [AFP]".

6. **Demonym**: One of the mentions is a demonym of the other. For example, "[Mexico] and [Mexican]". This depends on a list of demonyms extracted from the Wikipedia.

Although in the evaluation table 1 this sieve does not provide a huge increase in performance, it plays a crucial role in providing precise information with respect to the mention attributes that is later used in the Pronouns Match sieve.

3.4 Strict Head Match

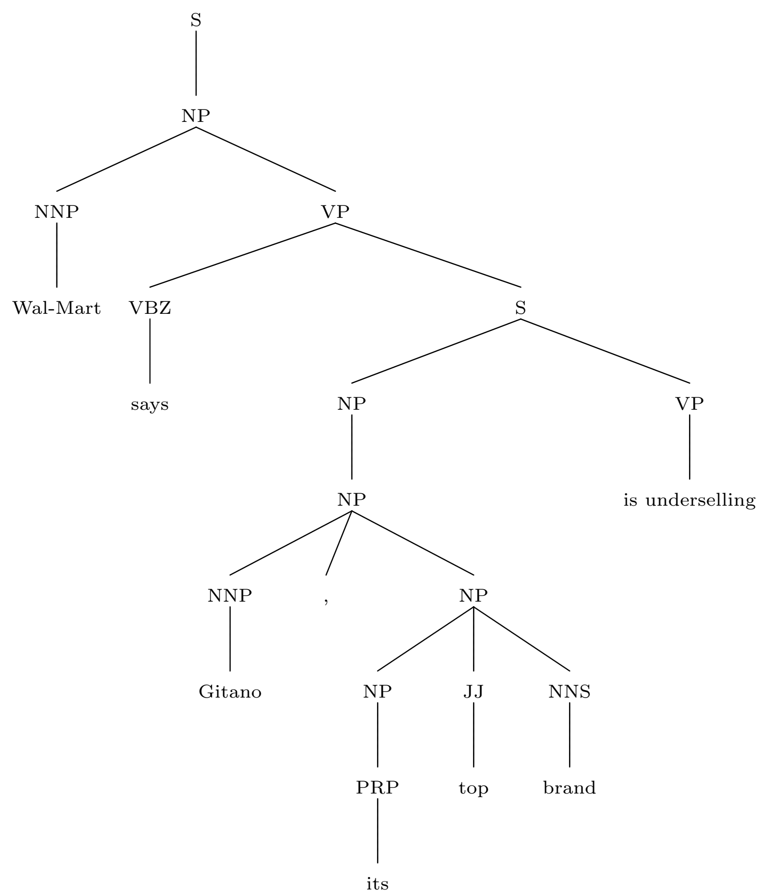

Figure 2: i-within-i construct.

Linking a mention to an antecedent based on the naive matching of their head words generates a lot of spurious links because it completely ignores possibly incompatible modifiers (Elsner and Charniak 2010). For example, *Yale University* and *Harvard University* have similar head words, but they are obviously different entities. To address this issue, this pass implements several constraints that **must all be matched** in order to yield a link:

1. **Entity head match**: The mention head word matches any head word of mentions in the antecedent entity. *This feature is constrained by enforcing a conjunction with the features below.*

2. **Word inclusion**: all the non-stop words in the current *entity* or mention cluster to be solved are included in the set of non-stop words in the *antecedent entity*. For example, this pass correctly clusters together the two mentions in the following text:

   (1) ... intervene in the [Florida Supreme Court]'s move ... does look like very dramatic change made by [the Florida court]

   and avoids clustering the two mentions in the following text:

   (2) The pilot had confirmed ... he had turned onto [the correct runway] but pilots behind him say he turned onto [the wrong runway].

3. **Compatible modifiers only**: the mention's modifiers, in this case only *nouns and adjectives*, are all included in the modifiers of the antecedent candidate. This feature models the same discourse property as the previous feature, but it focuses on the two individual mentions to be linked, rather than their corresponding entities or mention clusters.

4. **Not i-within-i**: the two mentions are not in an i-within-i construct. This means that one cannot be a child NP in the other's NP constituent. For example, in Figure 2 this prevents *its* from having the NP headed by *brand* in the set of possible antecedents. By propagation, it also removes the NP headed by *Gitano*. Therefore, it leaves the NP *Wal-Mart* as the closest compatible mention.

3.5 Pronominal Coreference Resolution

The implementation of this sieve pass is based on a well-known method, namely, enforcing constraints between the coreferent mentions. The following attributes are used for these constraints:

1. **Number**: Number attributes are assigned based on:

   (a) a static list for pronouns;

   (b) NER labels: mentions marked as a named entity are considered singular with the exception of organizations, which can be both singular and plural;

(c) part of speech tags: NN*S tags are plural and all other NN* tags are singular;

(d) a static dictionary from Bergsma and Lin (2006) [4].

2. **Gender**: gender attributes from static lexicons from Bergsma and Lin (2006) and from Ji and Lin (2009).

3. **Person**: Person attributes are assigned **only to pronouns**.

4. **Animacy**: Animacy attributes using:

   (a) a static list for pronouns;

   (b) NER labels, e.g., PERSON is animate whereas LOCATION is not;

   (c) a dictionary bootstrapped from the Web (Ji and Lin, 2009).

5. **NER label**.

6. **Pronoun distance**: sentence distance between a pronoun and its antecedent cannot be larger than 3.

EXCEPTION!!: The *person constraint for pronouns* is **NOT APPLIED** when linking two pronouns if one appears within quotes. This is a simple heuristic for speaker detection, e.g., I and she point to the same person in "[I] voted my conscience," [she] said.

Furthermore, it should be added that some languages provide genre and number information in their POS tags, as it is the Spanish case. In cases such as this, we first experiment with the information provided by the POS taggers instead of building static dictionaries for number and genre. This issue, however, does not affect animacy, for which the approach of Ji and Lin (2009).

4 Running Example

In this section, we will explain using an example how the mention processing as described in section 2.1 and section 2.2 is performed at each of the sieves which will be applied sequentially (Lee et al., 2013). We will describe each row of table 2 separately to obtain a clear view of the overall flow of the multi sieve pass coreference resolution system.

4.1 Mention Detection

1. In mention detection (section 2.1) the system extracts mentions by detecting NPs and other modifier pronouns (PRP). We identify 11 different mentions assigned them initially to 11 different mention clusters or entities. Remember that *superscript* marks cluster or entity id whereas *subscript* marks mention id.

[4]https://webdocs.cs.ualberta.ca/ bergsma/Gender/

Input	John is a musician. He played a new song. A girl was listening to the song. “It is my favorite,” John said to her.
Mention Detection	[John]$^{1}_{1}$ is [a musician]$^{2}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{6}_{6}$. “[It]$^{7}_{7}$ is [[my]$^{9}_{9}$ favourite]$^{8}_{8}$,” [John]$^{10}_{10}$ said to [her]$^{11}_{11}$.
Speaker Sieve	[John]$^{1}_{1}$ is [a musician]$^{2}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{6}_{6}$. “[It]$^{7}_{7}$ is [[**my**]$^{9}_{9}$ favourite]$^{8}_{8}$,” [**John**]$^{9}_{10}$ said to [her]$^{11}_{11}$.
String Match	[**John**]$^{1}_{1}$ is [a musician]$^{2}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{6}_{6}$. “[It]$^{7}_{7}$ is [[my]$^{9}_{9}$ favourite]$^{8}_{8}$,” [**John**]$^{1}_{10}$ said to [her]$^{11}_{11}$.
Precise Constructs	[**John**]$^{1}_{1}$ is [**a musician**]$^{1}_{2}$. [He]$^{3}_{3}$ played [a new song]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [the song]$^{6}_{6}$. “[**It**]$^{7}_{7}$ is [[**my**]$^{1}_{9}$ **favourite**]$^{7}_{8}$,” [John]$^{1}_{10}$ said to [her]$^{11}_{11}$.
Strict Head Match	[John]$^{1}_{1}$ is [a musician]$^{1}_{2}$. [He]$^{3}_{3}$ played [**a new song**]$^{4}_{4}$. [A girl]$^{5}_{5}$ was listening to [**the song**]$^{4}_{6}$. “[It]$^{7}_{7}$ is [[my]$^{1}_{9}$ favourite]$^{7}_{8}$,” [John]$^{1}_{10}$ said to [her]$^{11}_{11}$.
Pronoun Match	[**John**]$^{1}_{1}$ is [a musician]$^{1}_{2}$. [**He**]$^{1}_{3}$ played [a new song]$^{4}_{4}$. [**A girl**]$^{5}_{5}$ was listening to [**the song**]$^{4}_{6}$. “[**It**]$^{4}_{7}$ is [[my]$^{1}_{9}$ favourite]$^{7}_{8}$,” [John]$^{1}_{10}$ said to [**her**]$^{5}_{11}$.

Table 2: Running Coreference Resolution Example.

2. This step also extracts mention attributes from static lexicons referred to in section 3.5. For example, *John*{*ne:person*} and *A girl:*{*gender:female, number:singular*}.

The mention clusters which result after the application of the Mention Detection module form the input for the application of the sieve passes.

4.2 Speaker Detection Sieve

Although we do not provide an implementation of this sieve, we include this explanation for completeness reasons.

As explained in section 2.2.1, the selection of mentions (reducing search space) to be resolved is applied. In this case, as every mention cluster consists of one mention then every mention detected is a potential candidate for resolution (except the first mention in the text, obviously):

1. **Mention Selection**: Traverse every mention in the text from left to right and choose the first mention in every cluster.
2. The first match for candidate to be resolved by the Speaker Detection sieve pass is mention my_9^9 (because of the quotes).
3. The sieve pass links my_9^9 with $John_{10}^{10}$ into the same mention cluster or entity (cluster id 9).

4.3 Exact String Match Sieve

This sieve looks for antecedents that have the exact same string as the mention under consideration.

1. **Mention Selection**: Note that *this does not apply* in this sieve pass. So we also obtain $John_{10}^9$ as a candidate mention even though is is not the first mention in its cluster (my_9^9 is).
2. **Antecedent Selection**: As explained in section 2.2.2, candidate mentions are sorted from left to right using syntactic information. Thus, the resultant list of candidates for $John_{10}^9$ is: It_7^7, *My favorite*$_8^8$, My_9^9, *A girl*$_5^5$, *the song*$_6^6$, He_3^3, *a new song*$_4^4$, $John_1^1$, *a musician*$_2^2$.
3. **Exact String Match**: The algorithm stops when a matching antecedent is found. In this case, the algorithm will find $John_1^1$ and it stops there. Thus, $John_{10}^9$ changes mention cluster and becomes now $John_{10}^1$. The same goes to my_9^9: it now belongs to cluster 1 with $John_1^1$ and therefore becomes my_9^1.

4.4 Precise Constructs

The following row of table 2 shows the changes in the mention clusters when applying *Precise Constructs*. By applying this sieve pass *a musician*$_2^1$ is linked to $John_1^1$ and *my favorite*$_8^8$ is linked to It_7^7. In both cases the relevant construct is Subject-Predicate construction.

4.5 Strict Head Match

We do again the following steps:

1. **Mention Selection**: Textual order traversal of mentions to obtain first mention in their clusters as candidates. Among others, this provides *the song*$_6^6$ as candidate to be resolved.
2. **Antecedent Selection**: Left to right syntactic analysis procedure to obtain ordering of antecedents. This produces as the following order of antecedents for *the song*$_6^6$: *A girl*$_5^5$, He_3^3, *a new song*$_4^4$, and $John_1^1$.

3. **Strict Head Match**: Mentions without the same head word are removed. It only remains *a new* $song_4^4$.

4. **Strict Head Match**: non-stop words are all contained in antecedent and the 4 constraints of these sieve pass are held. Thus, *the* $song_6^6$ is becomes *the* $song_6^4$.

4.6 Pronoun Match

1. **Mention Selection**: We obtain as candidates for coreference linking He_3^3, It_7^7 and her_{11}^{11}.

2. **Antecedent Selection**: We obtain an antecedents ranking *for each of the candidates* obtained in the previous step. Thus, for He_3^3 we only get $John_1^1$ as antecedent. For It_7^7 we obtain *A* $girl_5^5$, He_3^3, *a new* $song_4^4$ and $John_1^1$.

3. **Pronoun match**: In the case of He_3^3 it is linked to $John_1^1$ cluster because of gender attribute. Gender attribute also helps to link her_{11}^{11} and *A* $girl_5^5$.

4. **Pronoun match**: It_7^7 is linked to *a new* $song_4^4$ based on the *animacy* attribute.

The multi-pass system approach also contemplates some post processing. For example, the singleton mention clusters are removed. They also removed, for the Ontonotes 4.0 corpus annotation the coreference links established in the *Precise Constructs* sieve shown in table 2.

5 Requirements

5.1 Linguistic Analyzers

1. **Constituent Parser**: con núcleos marcados.

2. **NERC**

3. **POS**: Morphological analysis (pronouns, etc.).

4. **Diccionaries**: demonyms, animacy, gender, number.

5.2 Datasets

1. **SemEval 2010**: datasets for 5 languages [5].

2. **CoNLL 2011 and 2012**: English coreference.

[5] `http://stel.ub.edu/semeval2010-coref/`

6 Exercises

(3) The disappearance of York University chef Claudia Lawrence is now being treated as suspected murder, North Yorkshire Police said. However detectives said they had not found any proof that the 35-year-old, who went missing on 18 March, was dead. Her father Peter Lawrence made a direct appeal to his daughter to contact him five weeks after she disappeared. His plea came at a news conference held shortly after a 10,000 reward was offered to help find Miss Lawrence. Crimestoppers said the sum they were offering was significantly higher than usual because of public interest in the case.

(Example 1 of RTE-4)

(4) American photojournalist James Nachtwey in a file photograph from May 18 2003 as he is awarded the Dan David prize in Tel Aviv for his out standing contribution to photography. It was announced by Time magazine on Thurs day, 11 December 2003 that Nachtwey was injured in Baghdad along with Time magazine senior correspondent Michael Weisskopf when a hand grenade was thrown into a Humvee they were traveling in with the US Army. Both journalists are reported in stable condition and are being evacuated to a US military hospital in Germany.

Caption 1470132 of the ImageCLEF-09 dataset (Paramita et al., 2009)

(5) Nothing special really. Comfortable and clean but very boring decor in comparison to other NH hotels. I stayed in NH in Brussels and Zurich and I really liked them because of their modern and stylish design and big rooms. This one was just like any other hotel. Basic rooms with basic and dull decor - bit disappointing. The customer service was average. The rate was very expensive and I still had to pay for Internet and 20 euros for breakfast!!! It was good but way overpriced! The best thing about the hotel was the location - city centre, 2min from a metro stop.

(review from OpeNER project)

References

Bergsma, S. and Lin, D. (2006). Bootstrapping path-based pronoun resolution. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, page 3340.

Haghighi, A. and Klein, D. (2009). Simple coreference resolution with rich syntactic and semantic features. *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing Volume 3 - EMNLP '09*, (August):1152.

Hobbs, J. R. (1978). Resolving pronoun references. *Lingua*, 44(4):311338.

Ji, H. and Lin, D. (2009). Gender and animacy knowledge discovery from web-scale n-grams for unsupervised person mention detection. *Proc. PACLIC2009*.

LDC, L. D. et al. (2005). *ACE (Automatic Content Extraction) English annotation guidelines for entities*. Version.

Lee, H., Chang, A., Peirsman, Y., Chambers, N., Surdeanu, M., and Jurafsky, D. (2013). Deterministic coreference resolution based on entity-centric, precision-ranked rules. *Computational Linguistics*, pages 1–54.

Lee, H., Peirsman, Y., Chang, A., Chambers, N., Surdeanu, M., and Jurafsky, D. (2011). Stanford's multi-pass sieve coreference resolution system at the CoNLL-2011 shared task. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning: Shared Task*, pages 28–34.

Paramita, M., Sanderson, M., and Clough, P. (2009). Diversity in photo retrieval: overview of the imageclefphoto task 2009. In *CLEF Working Notes 2009*.

Poon, H. and Domingos, P. (2008). Joint unsupervised coreference resolution with markov logic. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 650–659.

Raghunathan, K., Lee, H., Rangarajan, S., Chambers, N., Surdeanu, M., Jurafsky, D., and Manning, C. (2010). A multi-pass sieve for coreference resolution. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 492–501.

Weischedel, R., Pradhan, S., Ramshaw, L., Kaufman, J., Franchini, M., El-Bachouti, M., Xue, N., Palmer, M., Marcus, M., and Taylor, A. (2010). OntoNotes release 4.0. Technical report, Tech. rept. BBN Technologies.