Procesamiento del Lenguaje Natural, Revista nº 53, septiembre de 2014, pp 25-32 recibido 15-04-14 revisado 06-06-14 aceptado 10-06-14

Entity-Centric Coreference Resolution of Person Entities for Open Information Extraction *

Resolución de Correferencia Centrada en Entidades Persona para Extracción de Información Abierta

Marcos Garcia and **Pablo Gamallo**
Centro Singular de Investigación en Tecnoloxías da Información (CiTIUS)
Universidade de Santiago de Compostela
{marcos.garcia.gonzalez, pablo.gamallo}@usc.es

Resumen: Este trabajo presenta un sistema de resolución de correferencia de entidades *persona* cuya arquitectura se basa en la aplicación secuencial de módulos de resolución independientes y en una estrategia centrada en las entidades. Diversas evaluaciones indican que el sistema obtiene resultados prometedores en varios escenarios (≈ 71% y ≈ 81% de F1 CoNLL). Con el fin de analizar la influencia de la resolución de correferencia en la extracción de información, un sistema de extracción de información abierta se ha aplicado sobre textos con anotación correferencial. Los resultados de este experimento indican que la extracción de información mejora tanto en cobertura como en precisión. Las evaluaciones han sido realizadas en español, portugués y gallego, y todas las herramientas y recursos son distribuidos libremente.
Palabras clave: correferencia, anáfora, extracción de información abierta

Abstract: This work presents a coreference resolution system of person entities based on a multi-pass architecture which sequentially applies a set of independent modules, using an entity-centric approach. Several evaluations show that the system obtains promising results in different scenarios (≈ 71% and ≈ 81% F1 CoNLL). Furthermore, the impact of coreference resolution in information extraction was analyzed, by applying an open information extraction system after the coreference resolution tool. The results of this test indicate that information extraction gives better both recall and precision results. The evaluations were carried out in Spanish, Portuguese and Galician, and all the resources and tools are freely distributed.
Keywords: coreference, anaphora, open information extraction

1 Introduction

Relation Extraction (RE) systems automatically obtain different kinds of knowledge from unstructured texts, used for instance to build databases or to populate ontologies.

While RE systems usually depend on a set of predefined semantic relations for obtaining the knowledge (e.g., `hasHeadquartersAt`{*Organization, Location*}), Open Information Extraction (OIE) approaches perform unsupervised extractions of all types of verb-based relations (Banko et al., 2007).

In this respect, one of the main kind of extractions is related to person entities, and its objective is to obtain information about concrete people. For example, from a sentence such as "Obikwelu won the 100m gold medal at the 2009 Lusophony Games", an OIE system may obtain the following structured knowledge (with two arguments and a verb-based relation in each extraction):

OIE$_1$: *Obikwelu*$_{\text{Arg1}}$ `won` *the 100m gold medal*$_{\text{Arg2}}$

OIE$_2$: *Obikwelu*$_{\text{Arg1}}$ `won_the_100m_gold medal_at` *the 2009 Lusophony Games*$_{\text{Arg2}}$

However, many of the mentions of each person entity in a text are different from the others, so the final extraction may not be semantically complete. An OIE system could extract, from the same text than the previous example, relations like the following:

OIE$_3$: *Francis Obiorah Obikwelu*$_{\text{Arg1}}$ `is_a` *sprint athlete*$_{\text{Arg2}}$

OIE$_4$: *who*$_{\text{Arg1}}$ `is_based_in` *Lisbon*$_{\text{Arg2}}$

OIE$_5$: *He*$_{\text{Arg1}}$ `was_5`$^{\texttt{th}}$`_in` *the 200m*$_{\text{Arg2}}$

* This work has been supported by the HPCPLN project – Ref: EM13/041 (Galician Government) and by the Celtic – Ref: 2012-CE138 and Plastic – Ref: 2013-CE298 projects (Feder-Interconnecta).

ISSN 1135-5948

On the one hand, these extractions do not include the referents of the pronouns (*who*, *He*), so the knowledge could not be semantically useful. On the other hand, they do not report that *Obikwelu*, *Francis Obiorah Obikwelu* and *He* refer to the same entity, while *who* refers to other person.[1]

Related to that, Coreference Resolution (CR) systems use different techniques for clustering the various mentions of an entity into the same group. So, applying a coreference resolution tool before an OIE (or RE) system should improve the extraction in two ways: (1) increasing the recall by disambiguating the pronouns and (2) adding semantic knowledge by clustering both nominal and pronominal mentions of each entity.

This paper presents an open-source CR system for person entities which uses a multi-pass architecture. The approach, inspired by the Stanford Coreference Resolution System (Raghunathan et al., 2010), consists of a battery of modules applied from high-precision to high-recall. The system is also applied before a state-of-the-art OIE tool, in order to evaluate the impact of CR when performing information extraction.

The individual evaluations of the CR system show that the multi-pass architecture achieves promising performance when analyzing person entities ($\approx$ 71%/81% F1 CoNLL). Moreover, the OIE experiments prove that applying CR before an OIE system allows it to increase both the precision and the recall of the extraction.

Section 2 contains some related work and Section 3 presents the coreference resolution system. Then, Section 4 shows the results of both CR and OIE experiments while Section 5 points out the conclusions of this work.

2 Related Work

The different strategies for CR can be organized using two dichotomies: On the one hand, mention-pair *vs* entity-centric approaches. On the other hand, rule-based *vs* machine learning systems.

Mention-pair strategies decide if two mentions corefer using the features of these specific mentions, while entity-centric approaches take advantage of features obtained from other mentions of the same entities.

Rule-based strategies make use of sets of rules and heuristics for finding the best element to link each mention to (Lappin and Leass, 1994; Baldwin, 1997; Mitkov, 1998).

Machine learning systems rely on annotated data for learning preferences and constraints in order to classify pairs of mentions or entities (Soon, Ng, and Lim, 2001; Sapena, Padró, and Turmo, 2013). Some unsupervised models apply clustering approaches for solving coreference (Haghighi and Klein, 2007). Even though complex machine learning models obtain good results in this task, Raghunathan et al. (2010) presented a rule-based system that outperforms previous approaches. This system is based on a multi-pass strategy which first solves the *easy* cases, then increasing recall with further rules (Lee et al., 2013). Inspired by this method, *EasyFirst* uses annotated corpora in order to know whether coreference links are easier or harder (Stoyanov and Eisner, 2012).

For Spanish, Palomar et al. (2001) described a set of constraints and preferences for pronominal anaphora resolution, while Recasens and Hovy (2009) analyzed the impact of several features for CR. The availability of a large annotated corpus for Spanish (Recasens and Martí, 2010) also allowed other supervised systems being adapted for this language (Recasens et al., 2010).

Concerning OIE, several strategies were also applied since the first system, TextRunner (Banko et al., 2007). This tool (and further versions of it, such as ReVerb (Fader, Soderland, and Etzioni, 2011)) uses shallow syntax and labeled data for extracting triples (*argument_1, relation, argument_2*) which describe basic propositions.

Other OIE systems take advantage of dependency parsing for extracting the relations, such as WOE (Wu and Weld, 2010), which uses a learning-based model, or *DepOE* (Gamallo, Garcia, and Fernández-Lanza, 2012), a multi-lingual rule-based approach.

3 LinkPeople

This section describes LinkPeople, an entity-centric system which sequentially applies a battery of CR modules (Garcia and Gamallo, 2014a). Its architecture is inspired by Raghunathan et al. (2010), but it adds new modules for both cataphoric and elliptical pronouns as well as a set of syntactic constraints which

[1] In this paper, a *mention* is every instance of reference to a person, while an *entity* is all the mentions referring to the same person in the text (Recasens and Martí, 2010).

	Mention identif. ↓
Nominal Coreference	StringMatch NP_Cataphora PN_StMatch PN_Inclusion PN_Tokens HeadMatch Orphan_NP ↓
Pronominal Coreference	Pro_Cataphora Pronominal Pivot_Ent

Figure 1: Architecture of the system.

Who was $_1$[the singer of the Beatles]$_1$.
$_2$[The musician John Winston Ono Lennon]$_1$
was one of the founders of the Beatles.
With $_3$[Paul McCartney]$_2$, $_4$[he]$_1$ formed a
songwriting partnership. $_5$[Lennon]$_1$ was
born at Liverpool Hospital to $_6$[Julia]$_3$ and
$_7$[Alfred Lennon]$_4$. $_{8/9}$[$_{10}$[His]$_1$ parents]$_{3/4}$
named $_{11}$[him]$_1$ $_{12}$[John Winston Lennon]$_1$.
$_{13}$[Lennon]$_1$ revealed a rebellious nature
and acerbic wit. $_{14}$[The musician]$_1$ was
murdered in 1980.

Figure 2: Example of a coreference annotation of person entities. Mentions appear inside brackets. Numbers at the left are mention ids, while entity ids appear at the right.

restrict some links between pronominal and nominal mentions.

3.1 Architecture

The different modules of LinkPeople are applied starting from the most precise ones. Thus, *easy* links between mentions are done first. The entity-centric approach allows the system to use, in further modules, features that had been extracted in previous passes.

Figure 1 summarizes the architecture of LinkPeople. It starts with a mention identification module which selects the markables in a text. After that, a battery of nominal CR modules is executed. Finally, the pronominal solving passes —including a set of syntactic constraints— are applied.

Figure 2 contains a text with coreference annotation of person entities, used for exemplifying the architecture of LinkPeople.

In the first stage, a specific pass identifies the mentions referring to a person entity, using the information provided by the PoS-tagger and the NER as well as applying basic approaches for NP and elliptical pronoun identification: First, personal names (and noun phrases including personal names) are identified. Then, it seeks for NPs whose head may refer to a person (e.g., "the singer"). Finally, this module selects singular possessives and applies basic rules for identifying relative, personal and elliptical pronouns (Ferrández and Peral, 2000). At this step, each mention belongs to a different entity. Each entity contains the gender, number, head of a noun phrase, head of a Proper Noun (PN) and full proper noun as features. Once the mentions are identified, the CR modules are sequentially executed.

Each module applies the following strategy (except for some exceptional rules, explained below): mentions are traversed from the beginning of the text and each one is *selected* if (i) it is not the first mention of the text and (ii) it is the first mention of its entity. Once a mention is selected, it looks backwards for *candidates* in order to find an appropriate antecedent. If an antecedent is found, mentions are merged together in the same entity. Then, the next selected mention is evaluated. Apart from the mention identification pass, the current version of LinkPeople contains the following modules:

StringMatch: performs strict matching of the whole string of both mentions (the selected one and the candidate). In the example (Figure 2), mentions 13 and 5 are linked in this step.

NP_Cataphora: verifies if the first mention is a NP without a personal name. If so, it is considered a cataphoric mention, and the system checks if the next sentence contains a personal name as a Subject. In this case, these mentions are linked if they agree in gender and number. In the example, mentions 1 and 2 merge. Note that, at the end of this pass, this entity has as NP heads the words 'singer' and 'musician', and 'John Winston Ono Lennon' as the PN.

PN_StMatch: looks for mentions which share the whole PN, even if their heads are different (or if one of them does not have head). "The musician John Lennon" and "John Lennon" (not in Figure 2) would be an example.

PN_Inclusion: verifies if the full PN of the selected mention (in the entity) includes the

PN of the candidate mention (also in the entity), or vice-versa. In the example, this rule links mentions 5 and 2.

PN_Tokens: splits the full PN of a partial entity in its tokens, and verifies if the full PN of the candidate contains all the tokens in the same order, or vice-versa (except for some stop-words, such as "Sr.", "Jr.", etc.). As the pair "John Winston Ono Lennon" - "John Winston Lennon" are compatible, mentions 12 and 5 are merged.

HeadMatch: checks if the selected mention and the candidate one share the heads (or the heads of their entities). In Figure 2, mention 14 is linked to mention 13.

Orphan_NP: applies a pronominal-based rule to orphan NPs. A definite NP is marked as orphan if it is still a singleton and it does not contain a personal name. Thus, an orphan NP is linked to the previous PN with gender and number agreement. In the example, the mentions 8/9 are linked to 7 and 6.

Pro_Cataphora: verifies if a text starts with a personal (or elliptical) pronoun. If so, it seeks in the following sentence if there is a compatible PN.

Pronominal: this is the standard module for pronominal CR. For each selected pronoun, it verifies if the candidate nominal mentions satisfy the syntactic (and morphosyntactic) constraints. They include a set of constraints for each type of pronoun, which remove a candidate if any of them is violated. Some of them are: an object pronoun (direct or indirect) cannot corefer with its subject (mention 11 *vs* mentions 8/9); a personal pronoun does not corefer with a mention inside a prepositional phrase (mention 4 *vs* mention 3), a possessive cannot corefer with the NP it belongs to (mention 10 *vs* mentions 8/9) or a pronoun prefers a subject NP as its antecedent (mentions 10 and 11 *vs* mentions 6 and 7). This way, in Figure 2 the pronominal mention 4 is linked to mention 2, and mentions 10 and 11 to mention 5. This module has as a parameter the number of previous sentences for looking for candidates.

Pivot_Ent: this module is only applied if there are orphan pronouns (not linked to any proper noun/noun phrase) at this step. First, it verifies if the text has a pivot entity, which is the most frequent personal name in a text whose frequency is at least 33% higher than the second person with more occurrences. Then, if there is a pivot entity, all the orphan pronouns are linked to its mention. If not, each orphan pronoun is linked to the previous PN/NP (with no constraint).

4 Experiments

This section contains the performed evaluations. First, several experiments on CR are described. Then, a test of an OIE system is carried out, analyzing how LinkPeople influences the results of the extraction. All the experiments were performed in three languages: Spanish, Portuguese and Galician.[2]

4.1 Coreference Resolution

This section performs several tests with LinkPeople, comparing its results in different scenarios. First, the performance of LinkPeople using corpora with the mentions already identified (*gold-mentions*). Then, the basic mention identification module described in Section 3.1 is applied (*system-mentions*).[3] Both gold-mentions and system-mentions results were obtained with predicted information (*regular setting*): lemmas and PoS-tags (Padró and Stanilovsky, 2012; Garcia and Gamallo, 2010), NER (Padró and Stanilovsky, 2012; Garcia, Gayo, and González López, 2012; Gamallo and Garcia, 2011), and dependency annotation (Gamallo and González López, 2011), and with any kind of external knowledge (*closed setting*).

The experiments were performed with a Spanish corpus (46k tokens and $\approx 4,500$ mentions), a Portuguese one (51k tokens and $\approx 4,000$ mentions) and a Galician dataset (42k tokens and $\approx 3,500$ mentions) (Garcia and Gamallo, 2014b).

Three baselines were used: (i) *Singletons*, where every mention belongs to a different entity. (ii) *All_in_One*, where all the mentions belong to the same entity and (iii) *HeadMatch_Pro*, which clusters in the same entity those mentions sharing the head, and links each pronoun to the previous nominal mention with gender and number agreement.[4]

[2] All the tools and resources are freely available at `http://gramatica.usc.es/~marcos/LinkP.tbz2`

[3] Except for elliptical pronouns, where the gold-mentions were used for preserving the alignment, required for computing the results. Experiments in Section 4.2 simulate a real scenario.

[4] Due to language and format differences, other CR systems could not be used for comparison (Lee et al., 2013; Sapena, Padró, and Turmo, 2013).

Language	Model	MUC			B^3			$CEAF_e$			CoNLL
		R	P	F1	R	P	F1	R	P	F1	F1
Spanish	HeadMatch_Pro	78.2	90.7	84.0	35.3	81.2	49.2	72.9	51.5	60.4	64.5
	LinkPeople	84.1	94.1	**88.8**	62.9	84.8	**72.2**	83.4	71.0	**76.7**	**79.2**
Portuguese	HeadMatch_Pro	76.0	91.2	82.9	46.0	85.8	59.9	76.7	49.2	59.9	67.6
	LinkPeople	82.7	92.7	**87.4**	65.8	84.5	**74.0**	84.4	67.9	**75.2**	**78.9**
Galician	HeadMatch_Pro	81.9	89.8	85.7	44.1	83.6	57.7	70.0	53.5	60.6	68.0
	LinkPeople	89.0	94.6	**91.7**	72.9	88.4	**79.9**	87.6	76.6	**81.7**	**84.4**

Table 1: Results of LinkPeople (gold-mentions) compared to the best baseline (HeadMatch_Pro).

Language	Model	MUC			B^3			$CEAF_e$			CoNLL
		R	P	F1	R	P	F1	R	P	F1	F1
Spanish	Singletons	-	-	-	9.2	90.2	16.7	63.3	8.3	14.7	10.5
	All_In_One	77.5	78.8	78.1	69.0	44.2	53.9	5.7	73.5	10.5	47.5
	HeadMatch_Pro	68.2	81.5	74.2	31.4	68.1	43.0	62.6	52.0	56.8	58.0
	StringMatch	65.1	81.7	72.5	26.5	70.4	38.5	63.3	42.1	50.6	53.8
	PN_StMatch	66.3	81.8	73.2	27.2	69.4	39.0	63.1	45.0	52.6	54.9
	PN_Inclusion	69.6	82.8	75.6	34.2	68.2	45.5	64.6	55.0	59.4	60.2
	PN_Tokens	69.7	82.8	75.7	35.0	68.2	46.2	64.7	55.5	59.8	60.6
	Pronominal	70.2	82.3	75.8	35.9	67.3	46.8	64.6	59.4	61.9	61.5
	LinkPeople	72.8	85.3	**78.5**	50.4	70.7	**58.9**	73.5	68.3	**70.8**	**69.4**
Portuguese	Singletons	-	-	-	12.9	88.2	22.5	62.5	11.1	18.8	13.8
	All_In_One	75.6	73.1	74.3	67.1	37.7	48.3	9.5	61.3	16.4	46.3
	HeadMatch_Pro	65.4	79.4	71.7	41.5	68.9	51.8	69.2	54.2	60.8	61.4
	StringMatch	61.4	79.4	69.2	35.0	71.6	47.0	70.8	46.0	55.8	57.3
	PN_StMatch	63.7	79.9	70.9	36.6	70.8	48.3	71.6	50.7	59.4	59.5
	PN_Inclusion	67.9	81.1	73.9	45.5	69.7	55.1	74.0	61.9	67.4	65.5
	PN_Tokens	68.1	81.1	74.0	45.8	69.7	55.3	74.0	62.4	67.7	65.7
	Pronominal	69.0	81.2	74.6	47.9	69.2	56.6	74.0	65.4	69.4	66.9
	LinkPeople	69.9	82.0	**75.5**	55.8	69.2	**61.8**	76.6	68.8	**72.5**	**69.9**
Galician	Singletons	-	-	-	12.7	84.0	22.0	67.5	10.0	17.4	13.1
	All_In_One	83.1	71.1	76.6	75.9	40.9	53.2	7.4	67.1	13.3	47.7
	HeadMatch_Pro	72.4	73.7	73.0	38.0	62.3	47.2	61.4	52.5	56.6	59.0
	StringMatch	68.7	73.8	71.2	31.5	65.1	42.4	66.4	45.3	53.8	55.8
	PN_StMatch	70.7	74.0	72.3	34.4	64.1	44.8	66.6	50.4	57.4	58.1
	PN_Inclusion	74.6	75.2	74.9	42.8	63.0	50.9	66.7	60.1	63.2	63.0
	PN_Tokens	74.9	75.2	75.1	43.5	63.0	51.5	66.9	61.0	63.8	63.4
	Pronominal	75.0	75.0	75.0	44.2	62.7	51.8	67.0	62.4	64.6	63.8
	LinkPeople	78.5	78.5	**78.5**	65.0	67.0	**66.0**	78.6	73.4	**75.9**	**73.5**

Table 2: Results of LinkPeople (system-mentions) compared to the baselines.

The results were obtained using four metrics: MUC (Vilain et al., 1995), B^3 (Bagga and Baldwin, 1998), $CEAF_{entity}$ (Luo, 2005) and CoNLL (Pradhan et al., 2011). They were computed with the CoNLL 2011 scorer.

Table 1 contains the results of the best baseline and of LinkPeople using gold-mentions (for the full results of this scenario see Garcia and Gamallo (2014a)).

Table 2 includes the results of the three baselines and the performance values of LinkPeople using different modules added incrementally.[5] Central rows of Table 2 (StringMatch > Pronominal) include two baseline rules which classify mentions not covered by the active modules: (1) nominal mentions not analyzed are singletons and (2) pronouns are linked to the previous mention with number and gender agreement.

In every language and scenario, *HeadMatch_Pro* obtains good results, (as Recasens and Hovy (2010) shown), with ≈ 10% (F1 CoNLL) more than *All_in_One*.

The first module of LinkPeople (StringMatch) obtains lower results than the *HeadMatch_Pro* baseline, but with better pre-

[5]For spatial reasons, results of the modules with less performance improvements are omitted.

cision (except with the CEAF_e metric). After including more matching modules (NP_Cataphora and PN_StMatch), the results are closer to the best baseline, while the addition of PN_Inclusion and PN_Tokens modules allows the system to surpass it.

Then, HeadMatch, Orphan_NP and PRO_Cataphora slightly improve the performance of the system, while the pronominal resolution module notoriously increases the results in every evaluation and language. At this stage, LinkPeople obtains $\approx$ 76% and $\approx$ 64% (F1 CoNLL) in the gold-mentions and system-mentions scenarios, respectively.

Finally, one of the main contributions to the performance of LinkPeople is the combination of the Pronominal module with the Pivot_Ent one. This combination reduces the scope of the pronominal module, thus strengthening the impact of the syntactic constraints. Furthermore, Pivot_Ent looks for a prominent person entity in each text, and links the orphan pronouns to this entity.

The results of LinkPeople ($\approx$ 81% —gold-mentions— and $\approx$ 71% —system-mentions) show that this approach performs well for solving the coreference of person entities in different languages and text typologies.

4.2 Open Information Extraction

In order to measure the impact of LinkPeople in OIE, the most recent version of *DepOE*, was executed on the output of the CR tool.

LinkPeople was applied using a system-mentions approach, and without external resources. Apart from that, a basic elliptical pronoun module was included, which looks for elliptical pronouns in sentence-initial position, after adverbial phrases and after prepositional phrases. All the linguistic information was predicted by the same NLP tools referred in Section 4.1.

One corpus for each of the three languages was collected for performing the experiments. Each corpus contains 5 articles from Wikipedia and 5 from online journals.

DepOE was applied two times: First, using as input the plain text of the selected corpora (DepOE). Then, applied on the output of LinkPeople (DepOE+).

For computing precision of *DepOE*, 300 randomly selected triples containing at least a mention of a person entity as one of its arguments were manually revised (100 per language: 50 from Wikipedia and 50 from the journals). In the first run (without CR), extractions with pronouns as arguments were not computed, since they were considered as semantically underspecified. Thus, the larger number of extractions in the second run (DepOE+) is due to the identification of personal (including elliptical) pronouns. The central column of Table 3 contains an example of a new extraction obtained by virtue of CR.

LinkPeople also linked nominal mentions with different forms (right column in Table 3), thus enriching the extraction by allowing the OIE system to group various information of the same entity. An estimation of this improvement was computed as follows: from all the correct (revised) triples, it was verified if the personal mention in the argument had been correctly solved by LinkPeople. These cases were divided by the total number of correct triples, being these results considered as the *enrichment* value.

Table 4 contains the results of both DepOE and DepOE+ runs. DepOE+ was capable of extracting 22.7% more triples than the simple model, and its precision increased in about 10.6%. These results show that the improvement was higher in Wikipedia. This is due to the fact that the largest (person) entity in encyclopedic texts is larger than those in journal articles. Besides, Wikipedia pages contain more anaphoric pronouns referring to person entities (Garcia and Gamallo, 2014b). Finally, last column of Table 4 includes the percentage of enrichment of the extraction after the use of LinkPeople. Even tough these values are not a direct evaluation of OIE, they suggest that the information extracted by an OIE system is about 79% better when obtained after the use of a CR tool.

5 Conclusions

This paper presented LinkPeople, an entity-centric coreference resolution system for person entities which uses a multi-pass architecture and a set of linguistically motivated modules. It was evaluated in three languages, using different scenarios and evaluation metrics, achieving promising results.

The performance of the system was also evaluated in a real-case scenario, by analyzing the impact of coreference solving for open information extraction. The results show that using LinkPeople before the application of an OIE system allows to increase the performance of the extraction.

Sentence	"Debutó en la Tercera división"	"Anderson viajou por Europa"
DepOE	∅	*Anderson* `viajou_por` *Europa*
DepOE+	*Ander Herrera* `debutó_en` *la Tercera división*	*Wes Anderson* `viajou_por` *Europa*

Table 3: Extraction examples of DepOE and DepOE+ in Spanish (left) and Portuguese (right). The DepOE+ extraction in Spanish extracts a new triple —not obtained by DepOE— from a sentence with elliptical subject, while the first argument of the Portuguese example is enriched with the full proper name (and linked to other mentions in the same text).

Lg	**DepOE**			**DepOE+**			**E**
	W	**J**	**P**	**W**	**J**	**P**	
Sp.	47	82	49%	80	86	58%	84%
Pt.	82	133	39%	111	155	56%	75%
Gl.	168	114	49%	221	115	54%	77%

Table 4: Results of the two runs of *DepOE*. *W* and *J* include the number of extractions from Wikipedia and journalistic articles, respectively. *P* is the precision of the extraction, and *E* refers to the quality enrichment provided by LinkPeople.

In further work, the implementation of rules for handling plural mentions is planned, together with the improvement of nominal and pronominal constraints.

References

Bagga, Amit and Breck Baldwin. 1998. Algorithms for scoring coreference chains. In *Proceedings of the Workshop on Linguistic Coreference at the 1st International Conference on Language Resources and Evaluation*, volume 1, pages 563–566.

Baldwin, Breck. 1997. CogNIAC: high precision coreference with limited knowledge and linguistic resources. In *Proceedings of a Workshop on Operational Factors in Practical, Robust Anaphora Resolution for Unrestricted Texts*, pages 38–45.

Banko, Michele, Michael J Cafarella, Stephen Soderland, Matt Broadhead, and Oren Etzioni. 2007. Open information extraction from the web. In *Proceedings of the International Joint Conference on Artifical Intelligence*, pages 2670–2676.

Fader, Anthony, Stephen Soderland, and Oren Etzioni. 2011. Identifying relations for open information extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1535–1545.

Ferrández, Antonio and Jesús Peral. 2000. A computational approach to zero-pronouns in spanish. In *Proceedings of the Annual Meeting on Association for Computational Linguistics*, pages 166–172.

Gamallo, Pablo and Marcos Garcia. 2011. A resource-based method for named entity extraction and classification. In *Progress in Artificial Intelligence (LNCS/LNAI)*, volume 7026/2011, pages 610–623.

Gamallo, Pablo, Marcos Garcia, and Santiago Fernández-Lanza. 2012. Dependency-based Open Information Extraction. In *Proceedings of the Joint Workshop on Unsupervised and Semi-Supervised Learning in NLP*, pages 10–18.

Gamallo, Pablo and Isaac González López. 2011. A Grammatical Formalism Based on Patterns of Part-of-Speech Tags. *International Journal of Corpus Linguistics*, 16(1):45–71.

Garcia, Marcos and Pablo Gamallo. 2010. Análise Morfossintáctica para Português Europeu e Galego: Problemas, Soluções e Avaliação. *Linguamática*, 2(2):59–67.

Garcia, Marcos and Pablo Gamallo. 2014a. An Entity-Centric Coreference Resolution System for Person Entities with Rich Linguistic Information. In *Proceedings of the International Conference on Computational Linguistics*.

Garcia, Marcos and Pablo Gamallo. 2014b. Multilingual corpora with coreference annotation of person entities. In *Proceedings of the Language Resources and Evaluation Conference*, pages 3229–3233.

Garcia, Marcos, Iria Gayo, and Isaac González López. 2012. Identificação e Classificação de Entidades Mencionadas em Galego. *Estudos de Lingüística Galega*, 4:13–25.

Haghighi, Aria and Dan Klein. 2007. Unsupervised coreference resolution in a nonparametric bayesian model. In *Proceedings of the Annual Meeting on Association for Computational Linguistics*, volume 45, pages 848–855.

Lappin, Shalom and Herbert J. Leass. 1994. An algorithm for pronominal anaphora resolution. *Computational linguistics*, 20(4):535–561.

Lee, Heeyoung, Angel Chang, Yves Peirsman, N. Chambers, Mihai Surdeanu, and Dan Jurafsky. 2013. Deterministic coreference resolution based on entity-centric, precision-ranked rules. *Computational Linguistics*, 39(4):885–916.

Luo, Xiaoqiang. 2005. On Coreference Resolution Performance Metrics. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 25–32.

Mitkov, Ruslan. 1998. Robust pronoun resolution with limited knowledge. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics and International Conference on Computational Linguistics*, volume 2, pages 869–875.

Padró, Lluís and Evgeny Stanilovsky. 2012. FreeLing 3.0: Towards Wider Multilinguality. In *Proceedings of the Language Resources and Evaluation Conference*.

Palomar, Manuel, Antonio Ferrández, Lidia Moreno, Patricio Martínez-Barco, Jesús Peral, Maximiliano Saiz-Noeda, and Rafael Muñoz. 2001. An algorithm for anaphora resolution in Spanish texts. *Computational Linguistics*, 27(4):545–567.

Pradhan, Sameer, Lance Ramshaw, Mitchell Marcus, Martha Palmer, Ralph Weischedel, and Nianwen Xue. 2011. CoNLL-2011 Shared Task: Modeling Unrestricted Coreference in OntoNotes. In *Proceedings of the 15th Conference on Computational Natural Language Learning: Shared Task*, pages 1–27.

Raghunathan, Kathik, Heeyoung Lee, Sudarshan Rangarajan, Nathanael Chambers, Mihai Surdeanu, Dan Jurafsky, and Christopher Manning. 2010. A multi-pass sieve for coreference resolution. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 492–501.

Recasens, Marta and Eduard Hovy. 2009. A deeper look into features for coreference resolution. In *Anaphora Processing and Applications*. pages 29–42.

Recasens, Marta and Eduard Hovy. 2010. Coreference resolution across corpora: Languages, coding schemes, and preprocessing information. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 1423–1432.

Recasens, Marta and M. Antònia Martí. 2010. AnCora-CO: Coreferentially annotated corpora for Spanish and Catalan. *Language Resources and Evaluation*, 44.4:315–345.

Recasens, Marta, Lluís Màrquez, Emili Sapena, M. Antònia Martí, Mariona Taulé, Véronique Hoste, Massimo Poesio, and Yannick Versley. 2010. SemEval-2010 Task 1: Coreference resolution in multiple languages. In *Proceedings of the International Workshop on Semantic Evaluation*, pages 1–8.

Sapena, Emili, Lluís Padró, and Jordi Turmo. 2013. A Constraint-Based Hypergraph Partitioning Approach to Coreference Resolution. *Computational Linguistics*, 39(4).

Soon, Wee Meng, Hwee Tou Ng, and Daniel Chung Yong Lim. 2001. A machine learning approach to coreference resolution of noun phrases. *Computational linguistics*, 27(4):521–544.

Stoyanov, Veselin and Jason Eisner. 2012. Easy-first coreference resolution. In *Proceedings of the International Conference on Computational Linguistics*, pages 2519–2534.

Vilain, Marc, John Burger, John Aberdeen, Dennis Connolly, and Lynette Hirschman. 1995. A model-theoretic coreference scoring scheme. In *Proceedings of Message Understanding Conference 6*, pages 45–52.

Wu, Fei and Daniel S. Weld. 2010. Open information extraction using Wikipedia. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 118–127.