

Improving Machine Translation of Null Subjects in Italian and Spanish

Lorenza Russo, Sharid Loáiciga, Asheesh Gulati

Language Technology Laboratory (LATL)

Department of Linguistics – University of Geneva

2, rue de Candolle – CH-1211 Geneva 4 – Switzerland

{lorenza.russo, sharid.loaiciga, asheesh.gulati}@unige.ch

Abstract

Null subjects are non overtly expressed subject pronouns found in *pro-drop* languages such as Italian and Spanish. In this study we quantify and compare the occurrence of this phenomenon in these two languages. Next, we evaluate null subjects' translation into French, a “non pro-drop” language. We use the Europarl corpus to evaluate two MT systems on their performance regarding null subject translation: Its-2, a rule-based system developed at LATL, and a statistical system built using the Moses toolkit. Then we add a rule-based preprocessor and a statistical post-editor to the Its-2 translation pipeline. A second evaluation of the improved Its-2 system shows an average increase of 15.46% in correct pro-drop translations for Italian-French and 12.80% for Spanish-French.

1 Introduction

Romance languages are characterized by some morphological and syntactical similarities. Italian and Spanish, the two languages we are interested in here, share the *null subject parameter*, also called the *pro-drop parameter*, among other characteristics. The null subject parameter refers to whether the subject of a sentence is overtly expressed or not (Haegeman, 1994). In other words, due to their rich morphology, Italian and Spanish allow non lexically-realized subject pronouns (also called null subjects, zero pronouns or pro-drop).¹

From a monolingual point of view, regarding Spanish, previous work by Ferrández and Peral

¹Henceforth, the terms will be used indiscriminately.

(2000) has shown that 46% of verbs in their test corpus had their subjects omitted. Continuation of this work by Rello and Ilisei (2009) has found that in a corpus of 2,606 sentences, there were 1,042 sentences without overtly expressed pronouns, which represents an average of 0.54 null subjects per sentence. As for Italian, many analyses are available from a descriptive and theoretical perspective (Rizzi, 1986; Cardinaletti, 1994, among others), but to the best of our knowledge, there are no corpus studies about the extent this phenomenon has.²

Moreover, although null elements have been largely treated within the context of Anaphora Resolution (AR) (Mitkov, 2002; Le Nagard and Koehn, 2010), the problem of translating from and into a pro-drop language has only been dealt with indirectly within the specific context of MT, as Gojun (2010) points out.

We address three goals in this paper: i) to compare the occurrence of a same syntactic feature in Italian and Spanish, ii) to evaluate the translation of null subjects into French, that is, a “non pro-drop” language; and, iii) to improve the translation of null subjects in a rule-based MT system. Next sections follow the above scheme.

2 Null subjects in source corpora

We worked with the Europarl corpus (Koehn, 2005) in order to have a parallel comparative corpus for Italian and Spanish. From this corpus, we manually analyzed 1,000 sentences in both languages. From these 1,000 sentences (26,757 words for Italian, 27,971 words for Spanish), we identified 3,422 verbs for Italian and 3,184 for

²Poesio et al. (2004) and Rodríguez et al. (2010), for instance, focused on anaphora and deixis.

Spanish. We then counted the occurrences of verbs with pro-drop and classified them in two categories: personal pro-drop³ and impersonal pro-drop⁴, obtaining a total amount of 1,041 pro-drop in Italian and 1,312 in Spanish. Table 1 shows the results in percentage terms.

	Total Verbs	Pers. pro-drop	Impers. pro-drop	Total pro-drop
IT	3,422	18.41%	12.01%	30.42%
ES	3,182	23.33%	17.84%	41.17%

Table 1: Results obtained in the detection of pro-drop.

Results show a higher rate of pro-drop in Spanish (10.75%). It has 4.92% more personal pro-drop and 5.83% more impersonal pro-drop than Italian. The contrast of personal pro-drop is due to a syntactic difference between the two languages. In Spanish, sentences like (1a.) make use of two pro-drop pronouns while the same syntactic structure uses a pro-drop pronoun and an infinitive clause in Italian (1b.), hence, the presence of more personal pro-drop in Spanish.

- (1) a. ES *pro* le pido (1.sg) que *pro* intervenga (3.sg) con el prestigio de su cargo.
 b. IT *pro* le chiedo (1.sg) di intervenire (inf.) con il prestigio della sua carica.
I ask you to intervene with the prestige of your position.

The difference of impersonal pro-drop, on the other hand, is due to the Spanish use of an impersonal construction (2a.) with the “*se*” particle. Spanish follows the schema “*se* + finite verb + non-finite verb”; Italian follows the schema “finite verb + *essere* (to be) + past participle” (2b.). We considered this construction formed by one more verb in Italian than in Spanish as shown in examples (2a.) and (2b.). This also explains the difference in the total amount of verbs (Table 1).

- (2) a. ES Se podrá corregir.
 b. IT Potrà essere modificato.
It can be modified.

³Finite verbs with genuinely referential subjects (i.e. *I, you, s/he, we, they*).

⁴Finite verbs with non-referential subjects (i.e. *it*).

We found a total number of non-expressed pronouns in our corpus comparable to those obtained by Rodríguez et al. (2010) on the Live Memories corpus for Italian and by Recasens and Martí (2010) on the Spanish AnCora-Co corpus (Table 2). Note that in both of these studies, they were interested in co-reference links, hence they did not annotate impersonal pronouns, claiming they are rare. On the other hand, we took all the pro-drop pronouns into account, including impersonal ones.

Corpus	Language	Result
Our corpus	IT	3.89%
Live Memories	IT	4.5%
Our corpus	ES	4.69%
AnCora-Co	ES	6.36%

Table 2: Null-subjects in our corpus compared to Live Memories and AnCora-Co corpora. Percentages are calculated with respect to the total number of words.

3 Baseline machine translation of null subjects

The 1,000 sentences of our corpus were translated from both languages into French (IT→FR; ES→FR) in order to assess if personal pro-drop and impersonal pro-drop were correctly identified and translated. We tested two systems: Its-2 (Wehrli et al., 2009), a rule-based MT system developed at the LATL; and a statistical system built using the Moses toolkit out of the box (Koehn et al., 2007). The latter was trained on 55,000 sentence pairs from the Europarl corpus and tuned on 2,000 additional sentence pairs, and includes a 3-gram language model.

Tables 3 and 4 show percentages of correct, incorrect and missing translations of personal and impersonal null subjects calculated on the basis of the number of personal and impersonal pro-drop found in the corpus.

We considered the translation correct when the null pronoun is translated by an overt pronoun with the correct gender, person and number features in French; otherwise, we considered it incorrect. Missing translation refers to cases where the null pronoun is not generated at all in the target language.

We chose these criteria because they allow us to evaluate the single phenomenon of null subject

Its-2				
Pair	Pro-drop	Correct	Incorrect	Missing
IT→FR	personal	66.34%	3.49%	30.15%
	impersonal	16.78%	18.97%	64.23%
	<i>average</i>	<i>46.78%</i>	<i>9.6%</i>	<i>43.61%</i>
ES→FR	personal	55.79%	3.50%	40.70%
	impersonal	29.29%	11.40%	59.29%
	<i>average</i>	<i>44.28%</i>	<i>6.93%</i>	<i>48.78%</i>

Table 3: Percentages of correct, incorrect and missing translation of zero-pronouns obtained by Its-2. Average is calculated on the basis of total pro-drop in corpus.

Moses				
Pair	Pro-drop	Correct	Incorrect	Missing
IT→FR	personal	71.59%	1.11%	27.30%
	impersonal	44.76%	11.43%	43.79%
	<i>average</i>	<i>61%</i>	<i>5.18%</i>	<i>33.81%</i>
ES→FR	personal	72.64%	2.02%	25.34%
	impersonal	54.56%	2.45%	42.98%
	<i>average</i>	<i>64.78%</i>	<i>2.21%</i>	<i>33%</i>

Table 4: Percentages of correct, incorrect and missing translation of zero-pronouns obtained by Moses. Average is calculated on the basis of total pro-drop in corpus.

translation. BLEU and similar MT metrics compute scores over a text as a whole. For the same reason, human evaluation metrics based on adequacy and fluency were not suitable either (Koehn and Monz, 2006).

Moses generally outperforms Its-2 (Tables 3 and 4). Results for the two systems demonstrate that instances of personal pro-drop are better translated than impersonal pro-drop for the two languages. Since rates of missing pronoun translations are considerable, especially for Its-2, results also indicate that both systems have problems resolving non-expressed pronouns for their generation in French. A detailed description for each system follows.

3.1 Results – Its-2

Its-2 obtains better results for IT→FR personal pro-drop (66.34%) than for ES→FR (55.79%), but worse for impersonal pro-drop translation (16.78% and 29.29% respectively, Table 3).

For IT→FR translation in particular, Its-2 usually translates an Italian personal pro-drop with an overt pronoun of incorrect gender in French. In fact, it tends to translate a female personal pro-drop with a masculine overt pronoun. This prob-

lem is closely related with that of AR: as the system does not have any module for AR, it cannot detect the gender of the antecedent, rendering the correct translation infeasible.

The number of correct translation of impersonal pro-drop is very low for both pairs. ES→FR reached 29.29%, while IT→FR only 16.78% (Table 3). The reason for these percentages is a regular mistranslation or a missing translation. As for the mistranslation, in the case of Spanish, Its-2 translates the “*se*” pronoun in impersonal sentences by the French sequence *qui est-ce que* (*who*) (3). We attribute this behaviour to a lack of generation rules.

(3) ES Por consiguiente, mi grupo solicita que **se** suprima este punto del día.

FR Par conséquent, mon groupe sollicite qu’**on** supprime ce point de l’agenda.

ITS-2 * Par conséquent, mon groupe sollicite **qui est-ce que** supprime ce point du jour.

Therefore, my group ask to delete this point from the agenda.

With respects to missing pronouns in the target language (Table 3), percentages are quite high in both translation pairs (43.61% and 48.78% average missing pronouns respectively), especially with impersonal pro-drop. Let us take the example of ES→FR translation (59.29% missing pronouns): Its-2 never generates the French expletive pronouns “*ce, c’, ça*” (*it*) (4a.). For IT→FR translation (64.23% missing pronouns), it almost never generates the French pronoun “*il*” (*it*) for the impersonal 3rd person pro-drop pronoun in Italian (4b.).

However, if the system generates the pronoun, it is likely to be a first or a second person singular pronoun (“*je, tu*” – *I, you*) in French, increasing then the percentages of incorrectly translated impersonal pro-drop.

- (4) a. ES No es pedir demasiado.
FR Ce n’est pas trop demander.
ITS-2 * Pas est demander trop.
It is not too much to ask.
- b. IT È vero che [...].
FR Il est vrai que [...].
ITS-2 * Vrai que [...].
It is true that [...].

3.2 Results – Moses

Moses produces the highest percentage of correct translations for both personal and impersonal pro-drop, particularly in ES→FR (72.64% and 54.56% respectively, Table 4).

When translating personal pro-drop from Italian, sometimes the system generates infinitive forms instead of finite verbs (5).

- (5) IT Naturalmente accettiamo questo emendamento.
FR Bien sûr, nous acceptons cet amendement.
MOSES Bien sûr accepter cet amendement.
Of course, we accept this amendment.

When translating impersonal pro-drop from Italian, it performs worse (44.76%) because it tends not to generate the expletive pronoun (6).

Furthermore, for both source languages, Moses translates the impersonal pro-drop usually corre-

- (6) IT Vorrei, come mi è stato chiesto da alcuni colleghi, osservare un minuto di silenzio.

FR J’aimerais, comme il m’a été demandé par certains collègues, observer une minute de silence.

MOSES J’aimerais, comme m’a été demandé par certains collègues, observer une minute de silence.

I would like, as some colleagues asked me, to observe a minute of silence.

sponding to French pronoun “*on*” as the first plural personal pronoun (“*nous*” – *we*) (7a. and 7b.).

- (7) a. IT Io credo che si debba dare la precedenza alla sicurezza.
FR Je crois qu’**on** doit donner la priorité à la sécurité.
MOSES Je crois que **nous** devons donner la priorité à la sécurité.
I think that priority should be given to the safety.

- b. ES Como han demostrado los eventos recientes, queda mucho por hacer sobre este tema.

FR Comme l’ont montré les événements récents, **on** a encore beaucoup à faire sur ce thème.

MOSES Comme l’ont montré les événements récents, **nous** avons encore beaucoup à faire sur ce thème.

As it has been shown by recent events, there is much left to do on this subject.

4 Its-2 improvements

On the basis of this first evaluation, we tried to improve Its-2 pronoun generation when translating from Italian and Spanish. Two new components were added to the translation pipeline: a rule-based preprocessor, and a statistical post-editor. This section presents them in detail, along with the resources they rely on.

4.1 Rule-based preprocessing

Preprocessing of input data is a very common task in natural language processing. Statistical systems often benefit from linguistic preprocessing

to deal with rich morphology and long distance re-ordering issues (Sadat and Habash, 2006; Habash, 2007). In our case, the idea behind this first component is to help the translation process of a rule-based system by reducing the amount of zero pronouns in the source document⁵, ensuring that subject pronouns get properly transferred to the target language.

In order to assess the effect of this approach, we implemented a rule-based preprocessor taking as input a document in Italian or Spanish and returning as output the same document with dropped subject pronouns restored. It relies on two resources: a list of personal and impersonal verbs, and a part-of-speech tagging of the source document. We first present these two resources before describing the approach in more detail.

List of personal and impersonal verbs

This list simply contains surface forms of verbs. For our experiment, these forms were extracted from a subset of the Europarl corpus, where pro-drop verbs were manually annotated as taking a personal pronoun or an impersonal pronoun. This limits the coverage, but ensures domain-specific verb usage.

Part-of-speech tagging of the source document

Its-2, being a transfer-based system, relies on a parser to construct the syntactic structure of the source language and, from there, it transfers the syntactic structure onto the target language. Its-2 uses Fips (Wehrli, 2007), a multilingual parser also developed at LATL. Apart from the projection of syntactic structures, Fips produces part-of-speech tagging.

Outline of the approach

These are the steps followed by the preprocessor:

1. Read a part-of-speech tagged sentence.
2. Whenever a verb with no subject is encountered, check if it is a personal verb or an impersonal verb.

⁵In Italian and Spanish, even if a native speaker would not use subject pronouns in a given sentence, the same sentence with overtly expressed subject pronouns is grammatical. There might be a pragmatic difference, as pronouns are used in these languages when emphasis or contrast is desired (Haegeman, 1994).

3. If it is a personal verb, generate the appropriate pronoun before the verb (the masculine form is generated for the third person); if it is an impersonal verb, do not generate any pronoun.
4. Check the sentence for reordering according to syntactic rules of the target language (e.g. move the generated pronoun before a proclitic already preceding the verb).

An example of preprocessed sentences is given in Figure 1.

4.2 Statistical post-editing

Since the work of Simard et al. (2007a), statistical post-editing (SPE) has become a very popular technique in the domain of hybrid MT. The idea is to train a statistical phrase-based system in order to improve the output of a rule-based system. The statistical post-editing component is trained on a corpus comprising machine translated sentences on the source side (translations produced by the underlying rule-based system), and their corresponding reference translations on the target side. In a sense, SPE “translates” from imperfect target language to target language. Both quantitative and qualitative evaluations have shown that SPE can achieve significant improvements over the output of the underlying rule-based system (Simard et al., 2007b; Schwenk et al., 2009).

We decided to incorporate a post-editing component in order to assess if this approach can specifically address the issue of dropped subject pronouns. We first present the training corpus before describing the approach in more detail.

Training corpus

To train the translation model, we translated a subset of the Europarl corpus using Its-2. The translations were then aligned with corresponding reference translations, resulting in a parallel corpus for each language pair, composed of 976 sentences for IT→FR and 1,005 sentences for ES→FR. We opted for machine translations also on the target side, rather than human reference translations, in order to ascertain if a parallel corpus produced in such a way, with significantly lesser cost and time requirements, could be an effective alternative for specific natural language processing tasks.

Source in Italian	<i>pro</i> La ringrazio, onorevole Segni, <i>pro</i> lo farò volentieri.
Preprocessed	Io la ringrazio , onorevole Segni , io lo farò volentieri . <i>I thank you, Mr. Segni, I will do it willingly.</i>
Source in Spanish	Todo ello, de conformidad con los principios que <i>pro</i> siempre hemos apoyado.
Preprocessed	Todo ello, de conformidad con los principios que nosotros siempre hemos apoyado. <i>All this, according to the principles we have always supported.</i>

Figure 1: Output of the preprocessor: the pronoun in the first sentence is generated before the proclitic *lo*, and the pronoun in the second sentence is generated before the adverb “*siempre*” (*always*).

We reused the language model trained for the Moses experiment of Section 3.

Outline of the approach

We trained the SPE component using the Moses toolkit out of the box. With this setup, the final translation in French of a source sentence in Italian or Spanish can be obtained in two simple steps:

1. Translate a sentence using Its-2.
2. Give the translated sentence as input to the SPE component; the output of the SPE component is the final translation.

An example of post-edited translations is given in Figure 2.

4.3 Combination of preprocessing and post-editing

The two components described in the previous sections can be used separately, or combined together as the first and last elements of the same translation pipeline. With respect to the generation of pronouns, error analysis showed that the setup producing the best translations was indeed the combination of preprocessing and post-editing, which was therefore used in the full post-evaluation described in the next section. The example in Figure 3 illustrates progressive improvements (with respect to the generation of pronouns) achieved by using preprocessing and post-editing over the baseline Its-2 translation.

5 Post-evaluation

After adding the two components, we manually re-evaluated the translations of the same 1,000 sentences. Table 5 show percentages of correct, incorrect and missing translation of null subjects

in a comparative way: the first columns show percentages obtained by the baseline Its-2 system⁶ while the second columns show percentages obtained by the improved system.

Results show higher percentages of correct pro-drop translation for both language pairs, with an average increase of 15.46% for IT→FR, and 12.80% for ES→FR. Specifically, percentages of personal pro-drop translation for both pairs increased almost the same rate: 13.18% for IT→FR; 13.34% for ES→FR. It was not the case for impersonal pro-drop, where rates of the first pair augmented (18.98%), while the latter decreased (12.11%).

We explain this behaviour to a particular difficulty encountered when translating from Spanish, a language that largely prefers subjunctive mood clauses to other structures such as infinitive clauses. The problem arises because subjunctive tenses have a less distinctive morphology, with the same conjugation for the first and third person singular (9).

- (9) ES *pro* le pido (1.sg) que *pro* estudie (3.sg) un borrador de carta.
FR Je vous demande d’étudier un brouillon de lettre.
ITS-2 *Je (1.sg) demande que j’étudie (1.sg) un brouillon de charte.
I ask you to study a draft letter.

As a consequence of the improvement of personal pro-drop, only incorrect impersonal pro-drop translations decreased (Table 5). Indeed, if we consider personal pro-drop translation, we think that by means of restoring the pronouns for finite verbs, we also amplified the issue of AR. For instance, Italian uses the third singular person

⁶These percentages have already been discussed in section 3 and, in particular, in Table 3.

Its-2 translation	Vous invite à voter à faveur de l’amendement approuvé à l’unanimité [...].
Post-edited	Je vous invite à voter en faveur de l’amendement approuvé à l’ unanimité [...]. <i>I invite you to vote in favour of the amendment unanimously approved [...].</i>
Its-2 translation	Je madame Présidente, voudrais réclamer l’attention sur un cas [...].
Post-edited	Madame la Présidente, je voudrais réclamer l’attention sur un cas [...]. <i>Madam President, I would like to draw the attention on a case [...].</i>

Figure 2: Output of the post-editor: the pronoun in the first sentence is restored, and the pronoun in the second sentence is moved to the correct position.

Source in Italian	<i>pro</i> E’ la prima volta che <i>pro</i> intervengo in Plenaria e <i>pro</i> devo ammettere di essere molto emozionato [...].
Preprocessed	E’ la prima volta che io intervengo in Plenaria e io devo ammettere di essere molto emozionato [...].
Baseline	Qu’est la première fois qu’intervient dans *Plenaria et admettre de m’est beaucoup émus [...].
Translation	
-after preprocessing	Est la première fois que j ’interviens dans *Plenaria et j ’admettre d’est beaucoup émus [...].
-with post-editing	Qu ’est la première fois qu’ intervient en plénière ont et admettre de s ’est très émus [...].
-using both	C ’est la première fois que j ’ intervien en plénière ont et j ’admettre d’ est très émus [...]. <i>It is the first time that I speak in the Plenary session and I admit to being [...].</i>

Figure 3: Comparison of translation outputs: preprocessing leads to a better analysis of the sentence by Fips, as suggested by the presence of the pronouns “j” (*I*), absent in the baseline translation, and post-editing further restores successfully the missing impersonal pronoun “C” (*It*), whereas post-editing without preprocessing has no effect on pronoun generation.

		Baseline Its-2			Improved Its-2		
Pair	Pro-drop	Correct	Incorrect	Missing	Correct	Incorrect	Missing
IT→FR	personal	66.34%	3.49%	30.15%	79.52%	5.87%	14.60%
	impersonal	16.78%	18.97%	64.23%	35.76%	13.13%	51.09%
	average	46.78%	9.6%	43.61%	62.24%	8.73%	29%
ES→FR	personal	55.79%	3.50%	40.70%	69.13%	7.81%	23.04%
	impersonal	29.29%	11.40%	59.29%	41.40%	8.07%	50.52%
	average	44.28%	6.93%	48.78%	57.08%	7.92%	34.98%

Table 5: Percentages of correct, incorrect and missing translation of zero-pronouns. Results obtained by improved Its-2. Average is calculated on the basis of total pro-drop in corpus.

(“lei”) as a form of polite treatment, while French uses the second plural person (“vous”); however, Its-2 translates a restored pronoun for a finite verb in the third singular person as a finite verb in the third singular person in French too (8).

Gender discrepancies are an AR issue as well. For IT→FR, the problem of a female personal pro-drop translated by a masculine overt pronoun

in French still remains.

Finally, rates of missing pronouns also decreased. In this case, improvements are significant: we obtained a gain of 14.61% for the IT→FR pair and 13.8% for the ES→FR pair. Specifically, we obtained better improvements for personal pro-drop than for impersonal pro-drop. For the latter we think that rates decreased

- (8) **IT** Signora Presidente, mi permetta di parlare.
FR Madame la Présidente, **permettez-moi** de parler.
ITS-2 Madame la Présidente, **permet-moi** de parler.
Madam President, let me speak.

thanks only to the post-editing phase. Indeed, as both Italian and Spanish do not have any possible overt pronoun to be restored in the pre-processing phase, any improvement responds to changes made by the post-editor. On the other hand, improvements obtained for personal pro-drop translation confirm that the combination of the pre-processing and the post-editing together can be very advantageous, as already discussed in section 4.3.

6 Future work

As already mentioned, we have not found the solution to some problems yet.

First of all, we would like to include an AR module in our system. As it is a rule-based system, some problems as the subject pronoun mis-translation in subordinated sentences can be fixed by means of more specific rules and heuristics. Besides, an approach based on binding theory (Büring, 2005) could be effective as deep syntactic information is available, even though limited. For example, binding theory does not contain any formalization on gender, reason why a specific statistical component could be a more ideal option in order to tackle aspects such as masculine/feminine pronouns.

Secondly, an overt pronoun cannot be restored from a finite impersonal verb without making the sentence ungrammatical; therefore, our approach is not useful for treating impersonal sentences. As a consequence, we think that an annotation of the empty category, as done by Chung and Gildea (2010), could provide better results.

Also, in order to correctly render the meaning of a preprocessed sentence, we plan to mark restored subject pronouns in such a way that the information about their absence/presence in the original text is preserved as a feature in parsing and translation.

Finally, we would like to use a larger corpus to train the SPE component and compare the effects

of utilizing machine translations on the target side versus human reference translations. Besides, we would like to further explore variations on the plain SPE technique, for example, by injecting Moses translation of sentences being translated into the phrase-table of the post-editor (Chen and Eisele, 2010).

7 Conclusion

In this paper we measured and compared the occurrence of one syntactic feature – the null subject parameter – in Italian and Spanish. We also evaluated its translation into a “non pro-drop” language, that is, French, obtaining better results for personal pro-drop than for impersonal pro-drop, for both Its-2 and Moses, the two MT systems we tested.

We then improved the rule-based system using a rule-based preprocessor to restore pro-drop as overt pronouns and a statistical post-editor to correct the translation. Results obtained from the second evaluation showed an improvement in the translation of both sorts of pronouns. In particular, the system now generates more pronouns in French than before, confirming the advantage of using a combination of preprocessing and post-editing with rule-based machine translation.

Acknowledgments

This work has been supported in part by the Swiss National Science Foundation (grant No 100015-130634).

References

- Daniel Büring. 2005. *Binding Theory*. Cambridge Textbooks in Linguistics.
- Anna Cardinaletti. 1994. Subject Position. *GenGenP*, 2(1):64–78.
- Yu Chen and Andreas Eisele. 2010. Hierarchical Hybrid Translation between English and German. In *Proceedings of the 14th Annual Conference of the European Association for Machine Translation*, pages 90–97.
- Tagyoung Chung and Daniel Gildea. 2010. Effects of Empty Categories on Machine Translation. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, EMNLP ’10, pages 636–645.
- Antonio Ferrández and Jesús Peral. 2000. A Computational Approach to Zero-pronouns in Spanish. In *Proceedings of the 38th Annual Meeting of the*

- Association for Computational Linguistics*, pages 166–172.
- Anita Gojun. 2010. Null Subjects in Statistical Machine Translation: A Case Study on Aligning English and Italian Verb Phrases with Pronominal subjects. Diplomarbeit, Institut für Maschinelle Sprachverarbeitung, University of Stuttgart.
- Nizar Habash. 2007. Syntactic Preprocessing for Statistical Machine Translation. In *Proceedings of Machine Translation Summit XI*, pages 215–222.
- Liliane Haegeman. 1994. *Introduction to Government and Binding Theory*. Blackwell Publishers.
- Philipp Koehn and Christof Monz. 2006. Manual and Automatic Evaluation of Machine Translation between European Languages. In *Proceedings on the HTL-NAACL Workshop on Statistical Machine Translation*, pages 102–121. Association for Computational Linguistics.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoli, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Christopher J. Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open Source Toolkit for Statistical Machine Translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180. Association for Computational Linguistics.
- Philipp Koehn. 2005. Europarl: A Parallel Corpus for Statistical Machine Translation. In *Proceedings of the 10th Machine Translation Summit (MT Summit X)*.
- Ronan Le Nagard and Philipp Koehn. 2010. Aiding Pronoun Translation with Co-Reference Resolution. In *Proceedings of the Joint 5th Workshop on Statistical Machine Translation*, pages 258–267.
- Ruslan Mitkov. 2002. *Anaphora Resolution*. Longman.
- Massimo Poesio, Rodolfo Delmonte, Antonella Bristot, Luminita Chiran, and Sara Ronelli. 2004. The VENEX corpus of anaphora and deixis in spoken and written Italian. In *Manuscript*.
- Marta Recasens and M. Antònia Martí. 2010. AnCora-CO: Coreferentially Annotated Corpora for Spanish and Catalan. *Language Resources and Evaluation*, 44(4):315–345.
- Luz Rello and Iustina Ilisei. 2009. A Rule-Based Approach to the Identification of Spanish Zero Pronouns. In *Proceedings of Student Research Workshop, RANLP*, pages 60–65.
- Luigi Rizzi. 1986. Null Objects in Italian and the Theory of *pro*. *Linguistic Inquiry*, 17(3):501–557.
- Kepa J. Rodríguez, Francesca Delogu, Yannick Versley, Egon W. Stemle, and Massimo Poesio. 2010. Anaphoric Annotation of Wikipedia and Blogs in the Live Memories Corpus. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC)*, Valletta, Malta. European Language Resources Association (ELRA).
- Fatiha Sadat and Nizar Habash. 2006. Combination of Arabic Preprocessing Schemes for Statistical Machine Translation. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the ACL*, pages 1–8.
- Holger Schwenk, Sadaf Abdul-Rauf, Loïc Barrault, and Jean Senellart. 2009. SMT and SPE Machine Translation Systems for WMT’09. In *Proceedings of the 4th Workshop on Statistical Machine Translation*, pages 130–134.
- Michel Simard, Cyril Goutte, and Pierre Isabelle. 2007a. Statistical Phrase-based Post-editing. In *Proceedings of NAACL HLT 2007*, pages 508–515.
- Michel Simard, Nicola Ueffing, Pierre Isabelle, and Roland Kuhn. 2007b. Rule-based Translation With Statistical Phrase-based Post-editing. In *Proceedings of the 2nd Workshop on Statistical Machine Translation*, pages 203–206, June.
- Eric Wehrli, Luka Nerima, and Yves Scherrer. 2009. Deep Linguistic Multilingual Translation and Bilingual Dictionaries. In *Proceedings of the Fourth Workshop on Statistical Machine Translation*, pages 90–94.
- Eric Wehrli. 2007. Fips, a “Deep” Linguistic Multilingual Parser. In *Proceedings of the Workshop on Deep Linguistic Processing*, pages 120–127. Association for Computational Linguistics.