

Italian and Spanish Null Subjects: A Case Study Evaluation in an MT Perspective

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

Thanks to their rich morphology, Italian and Spanish allow pro-drop pronouns, i.e., non lexically-realized subject pronouns. Here we distinguish between two different types of null subjects: personal pro-drop and impersonal pro-drop. We evaluate the translation of these two categories into French, a non pro-drop language, using Its-2, a transfer-based system developed at our laboratory; and Moses, a statistical system. Three different corpora are used: two subsets of the Europarl corpus and a third corpus built using newspaper articles. Null subjects turn out to be quantitatively important in all three corpora, but their distribution varies depending on the language and the text genre though. From a MT perspective, translation results are determined by the type of pro-drop and the pair of languages involved. Impersonal pro-drop is harder to translate than personal pro-drop, especially for the translation from Italian into French, and a significant portion of incorrect translations consists of missing pronouns.

Keywords: null subject, corpus study, MT evaluation.

1. Introduction

Null subjects are non overtly expressed subject pronouns found in *pro-drop* languages such as Italian and Spanish (Haegeman, 1994). They are also known as elliptical pronouns (Recasens and Martí, 2010), zero pronouns (Ferrández and Peral, 2000; Mitkov, 2002; Rello and Ilisei, 2009) and pro-drop pronouns (Rizzi, 1986). These terms will be used interchangeably.

Italian and Spanish allow non lexically-realized subject pronouns thanks to their rich morphology. In this article we distinguish between two different types of null subjects: personal pro-drop and impersonal pro-drop. In the first case, finite verbs have a genuinely referential non expressed subject; in the latter, they do not have a referential subject.

Corpus studies have treated pro-drop pronouns within the Anaphora Resolution (AR) context (Recasens and Martí, 2010; Rodríguez et al., 2010, among others). However, these studies do not consider impersonal pro-drop pronouns because they are not genuinely referential subjects, and therefore do not allow co-reference with an antecedent.

Few studies have been published within the context of Machine Translation (MT). Chung and Gildea (2010), for instance, do not distinguish between personal and impersonal pro-drop in their work on Chinese and Korean. Gojun (2010) mentions the distinction, but describes the results of her system in general terms using BLEU.

Aiming at a better understanding of null subjects and their translation into French, a non pro-drop language (IT→FR, ES→FR), we conducted an evaluation on two data sets extracted from the Europarl corpus (Koehn, 2005): the first, containing the same sentences in Italian and Spanish¹; the second subset, containing texts produced by native speak-

ers only.² We considered this second corpus as *comparable* under the definition given by McEnery and Xiao (2008) and it is intended to avoid the effect of *translationese*, as explained by Ilisei and Mihăilă (2009). Furthermore, we conducted a third evaluation on newspaper articles, in order to better understand the systems' performance on null subjects using raw texts³.

This article is organized as follows: we first evaluate the occurrence of pro-drop in Italian and Spanish in the two corpora (section 2). Then we translate these two corpora using two machine translation systems: Its-2, a transfer-based system developed in our laboratory; and Moses, a statistical system (section 3). We discuss our results in section 4. Finally, we describe a third evaluation on a news corpus (section 5). A general discussion is presented in section 6. Section 7 concludes the article.

2. Null Subjects in Corpora

An important feature of the Europarl corpus is that it is a parallel corpus. Paraphrasing McEnery and Xiao (2008), parallelism gives a good basis for studying how a syntactic phenomenon is conveyed into another language. However, this feature implies that at least one of the languages involved is a translation itself of another one. Therefore, in order to avoid the undeniable effect of *translationese* (which results in more explicit texts with less pro-drop pronouns, as it has been demonstrated by Ilisei and Mihăilă (2009)), we also conduct an evaluation using a native corpus, ensuring the completeness of the evaluation.

2.1. Null Subjects in Parallel Corpus

We worked with the Europarl corpus, release v3⁴, in order to have the same corpus for Italian and Spanish, as reported

¹We will refer to this corpus as “parallel corpus”.

²We will refer to this second corpus as “native corpus”.

³We will refer to this corpus as “news corpus”.

⁴<http://opus.lingfil.uu.se/Europarl13.php>

in Russo et al. (2012).

From this corpus, we manually analyzed 1 000 sentences in each language (26 757 words in Italian; 27 971 words in Spanish), and we identified 3 422 verbs in Italian and 3 184 in Spanish. We then counted the instances of verbs with pro-drop and classified them in two categories: personal pro-drop and impersonal pro-drop. We obtained a total amount of 1 041 pro-drop in Italian and 1 312 in Spanish. Table 1 shows this distribution in percentage.

	Italian	Spanish
personal pro-drop	18.41%	23.33%
impersonal pro-drop	12.01%	17.84%
total pro-drop	30.42%	41.17%

Table 1: Distribution of zero-pronouns in the Europarl parallel corpus.

The total rate of Spanish pro-drop pronouns is higher than the Italian one ($\chi^2(1, N = 6\,606) = 83.66, p < .05$), but there is no significant difference between the two languages with respect to the type of pronouns ($\chi^2(1, N = 2\,352) = 3.53, n.s.$).

2.2. Null Subjects in Native Corpus

We decided to build a comparable corpus exclusively composed of texts produced by native speakers, in order to confirm or contrast results obtained in the initial evaluation. We considered the same number of sentences (1 000) from a subset of the same corpus (Europarl, release v5 for Italian⁵, and release v6 for Spanish⁶).

The genre and the domain of the texts are the same as in our first data set, the only difference being that this second subset is exclusively composed of texts produced by native speakers (there are no sentences shared between the two data sets). We extracted a sample of 1 000 sentences from Europarl that were annotated with the ISO 639-1 code of the language the original speaker was using, either in the LANGUAGE attribute of the SPEAKER tag, or in parentheses at the start of a speech, in our case (IT) and (ES). The number of such annotated sentences was sufficient for Spanish, but not for Italian, we thus extended the coverage for Italian by including speeches by speakers from Italy (using the NAME attribute of the SPEAKER tag and checking the country of origin on the European Parliament/MEPs search engine⁷). We then obtained the same amount of sentences for both languages.

We manually evaluated this sample of 1 000 sentences (30 259 words in Italian, 34 442 in Spanish), finding 3 927 verbs in Italian and 3 844 in Spanish. We then counted the instances of personal and impersonal pro-drop verbs, obtaining a total amount of 1 285 pro-drop verbs in Italian and 1 371 in Spanish. Table 2 shows this distribution in percentage.

	Italian	Spanish
personal pro-drop	20.24%	25.99%
impersonal pro-drop	12.48%	9.68%
total pro-drop	32.72%	35.67%

Table 2: Distribution of zero-pronouns in the Europarl native corpus.

As in the parallel corpus, Spanish has more pro-drop pronouns than Italian ($\chi^2(1, N = 7\,771) = 7.48, p < .05$). But there is a significant difference between the two languages: Spanish has more personal pro-drop than Italian, and Italian has more impersonal pro-drop than Spanish ($\chi^2(1, N = 2\,656) = 36.60, p < .05$).

The smaller number of Spanish impersonal pro-drop comes from the increase of personal pro-drop. Indeed, we noticed the use of more 1st person pronouns (1a) and more referential noun phrases (1b).

- (1) a. ES *pro*Sigo el caso del Tíbet; *pro*he visitado la región; *pro*he hablado con muchas personas y *pro*considero legítimas sus reivindicaciones.
FR **Je** suis le cas du Tibet; **j'**ai visité la région; **j'**ai parlé avec beaucoup de personnes et **je** considère légitimes ses revendications.
EN *I follow the Tibet case; I have visit the region; I have spoken with many people and I have considered their vindications legitimate.*
- b. ES ¡Menos mal que el [señor Comisario]_i tenía pocas enmiendas que *pro*no [aceptara]_i.
FR Heureusement que Monsieur le Commissaire avait peu de modifications qu'**il** n'accepterait pas!
EN *Luckily Mr. Commissioner had few amendments not to accept!*

3. Machine Translation Evaluation

We tested two systems: Its-2 (Wehrli et al., 2009), a transfer-based MT system developed in our laboratory (LATL); and a statistical system built using the Moses Toolkit out of the box (Koehn et al., 2007). For both language pairs (IT→FR and ES→FR), Moses was trained using 55 000 sentence pairs and tuned on 2 000 sentence pairs. It includes a 3-gram language model.

The first columns of Tables 3 and 4 show results obtained on the parallel corpus, while the last columns show results obtained on the native corpus.⁸

The translation is considered correct when the null pronoun in the source language is translated as an overt personal pronoun with the correct gender, person and number features in French; otherwise, we considered it incorrect. Missing translation occurs when the null pronoun is not generated at all in the target language.⁹

⁵<http://www.statmt.org/europarl/archives.html#v5>

⁶<http://www.statmt.org/europarl/>

⁷<http://www.europarl.europa.eu/meps/>

⁸Results obtained on the parallel corpus have already been reported and discussed in Russo et al. (2012).

⁹ χ^2 calculations are done merging the incorrect and the missing categories.

		Its-2					
		Parallel Corpus			Native Corpus		
Pair	Pro-drop	Correct	Incorrect	Missing	Correct	Incorrect	Missing
IT→FR	personal	66.34%	3.49%	30.15%	67.42%	2.64%	29.93%
	impersonal	16.78%	18.97%	64.23%	13.46%	19.59%	66.93%
	average	46.78%	9.6%	43.61%	46.84%	9.10%	44.04%
ES→FR	personal	55.79%	3.50%	40.70%	46.65%	1.50%	51.85%
	impersonal	29.29%	11.40%	59.29%	34.95%	6.99%	58.06%
	average	44.28%	6.93%	48.78%	43.47%	2.99%	53.54%

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

		Moses					
		Parallel Corpus			Native Corpus		
Pair	Pro-drop	Correct	Incorrect	Missing	Correct	Incorrect	Missing
IT→FR	personal	71.59%	1.1%	27.30%	74.96%	1.13%	23.89%
	impersonal	44.76%	11.43%	43.79%	43.46%	10.00%	46.53%
	average	61.00%	5.18%	33.81%	62.95%	4.51%	32.52%
ES→FR	personal	72.64%	2.02%	25.34%	77.68%	0.70%	21.62%
	impersonal	54.56%	2.45%	42.98%	53.49%	2.69%	43.89%
	average	64.78%	2.21%	33.00%	71.12%	1.24%	27.64%

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

Translation results follow the same tendency across the corpora: personal pro-drop are better translated than impersonal pro-drop by both systems. (For the parallel corpus: Its-2 obtained the following results: IT→FR ($\chi^2(1, N = 1041) = 245.40, p < .05$); ES→FR ($\chi^2(1, N = 1312) = 91.73, p < .05$). Moses obtained the following results: IT→FR ($\chi^2(1, N = 1041) = 75.20, p < .05$); ES→FR ($\chi^2(1, N = 1312) = 30.59, p < .05$). For the native corpus: Its-2 obtained the following results: IT→FR ($\chi^2(1, N = 1285) = 354.38, p < .05$); ES→FR ($\chi^2(1, N = 1371) = 15.10, p < .05$). Moses obtained the following results: IT→FR ($\chi^2(1, N = 1285) = 128.98, p < .05$); ES→FR ($\chi^2(1, N = 1371) = 77.17, p < .05$).

Both systems make the same translation mistakes on both corpora. For instance, Its-2 does not distinguish between the 1st person singular “*je*” (I) and 3rd person singular “*il*” (s/he) pronouns for ES→FR (2a), because they both have the same verbal endings.

When translating from Italian into French, Its-2 almost never generates the French pronoun “*il*” for the 3rd person impersonal pro-drop pronoun in Italian (2b).

- (2) a. ES Se produjo el secuestro de un buque atunero cuando *pro*faenaba en aguas internacionales.
FR Il y a eu l’enlèvement d’un bateau pour la pêche au thon quand **il** était dans les eaux internationales.
ITS-2 Il n’a produit l’enlèvement d’un navire *atunero quand **je** *faenaba dans eaux internationales.
EN *There was the kidnapping of a tuna boat*

when fishing in international waters.

- b. IT È importante sottolineare questa differenza.
FR Il est important de souligner cette différence.
ITS-2 Est important de souligner cette différence.
EN *It is important to stress this difference.*

On the other hand, Moses often translates the 2nd person of polite treatment as the 3rd person plural when translating from Spanish into French (3a). Besides, for IT→FR, it always generates the 1st person plural pronoun “*nous*” (*we*) instead of the impersonal construction “*il y a*” (*there is*) (3b).

- (3) a. ES Para ustedes el mercurio es apenas un metal tóxico que solo *pro*han visto en los termómetros cuando *pro*tenían fiebre.
FR Pour vous le mercure est juste un métal toxique que **vous** n’avez vu que dans les thermomètres lorsque **vous** aviez la fièvre .
MOSES *Le mercure pour vous est à peine dégagés métal toxique seulement n’ont vu dans les *termomètres quand **ils** avaient fièvre.
EN *For you mercury is just a toxic metal you have only seen in thermometers when you had fever.*
- b. IT Ci sarà molto da fare.
FR **Il** y aura beaucoup à faire.
MOSES **Nous** aurons beaucoup à faire.

4. Discussion

Under the assumption that the parallel corpus is biased by *translationese*, we expected a difference in the occurrence of null subjects in the native corpus. In contrast, we expected to see the same average translation quality by both systems (since both corpora are subsets of the Europarl corpus).

There is indeed a difference between the parallel and the native corpora for Italian: native has slightly more pro-drop verbs than parallel ($\chi^2(1, N = 7\,349) = 4.48, p < .05$). By contrast, parallel has more pro-drop verbs than native in Spanish ($\chi^2(1, N = 7\,028) = 22.65, p < .05$). This confirms Ilisei and Mihăilă (2009)’s results only partially. We claim so, because –as in their study– we do have more personal pro-drop in the native corpus, since it is composed of texts produced by native speakers only (Tables 1 and 2). However, they do not consider impersonal pro-drop, whereas we do. This is why almost all the percentages of personal and impersonal pro-drop in the native corpus are higher when compared to the parallel corpus, except for Spanish impersonal pro-drop (as showed in examples 1a and 1b).

With regard to translation quality, both systems achieve the same performance on both corpora (the difference for correct pro-drop translation across both corpora is not significant (IT→FR ($\chi^2(1, N = 2\,533) = 0.14, n.s.$); for ES→FR ($\chi^2(1, N = 3\,002) = 2.23, n.s.$)). The systems’ performance decreases when translating impersonal pro-drop into French, especially for IT→FR. (For correct pro-drop translation: for IT→FR: ($\chi^2(1, N = 2\,533) = 85.27, p < .05$); for ES→FR ($\chi^2(1, N = 3\,002) = 2.68, n.s.$)). Thus, it seems necessary to further study this specific phenomenon.

5. Translation of Newspaper Articles

In order to further evaluate the systems’ performance, we also compared results obtained on the parallel corpus and the native corpus to another stylistically different data set.

5.1. News Corpus

For this third evaluation, we annotated 1 000 sentences (19 649 words) from the economics section of the Italian newspaper *La Repubblica*¹⁰, and 1 000 sentences (34 510 words) from the economics section of the Spanish newspaper *El Mundo*¹¹. Articles date from 2000.

From these 1 000 sentences, we counted 2 102 verbs in Italian and 2 932 verbs in Spanish, of which 453 have a pro-drop pronoun in Italian and 427 in Spanish (Table 5).

We found a different tendency in this data set: Italian has more pro-drop pronouns ($\chi^2(1, N = 5\,034) = 41.44, p < .05$) than Spanish. In particular, Italian displays a higher percentage for impersonal pro-drop ($\chi^2(1, N = 880) = 80.55, p < .05$) than for personal pro-drop. The reason is a massive use of the impersonal construction with the “*si*” clitic pronoun in Italian (4).

	Italian	Spanish
personal pro-drop	7.66%	9.58%
impersonal pro-drop	13.89%	4.98%
total pro-drop	21.55%	14.56%

Table 5: Distribution of zero-pronouns in the news corpus.

- (4) a. IT Si deve ridurre la spesa pubblica.
FR On doit réduire la dépense publique.
EN *We have to reduce the public expense.*

5.2. News Corpus Evaluation

For this task, we tested the same systems as before. This time, performance decreased for both systems (Table 6). Its-2 performs worse when translating impersonal pro-drop than personal pro-drop (IT→FR ($\chi^2(1, N = 453) = 54.88, p < .05$); ES→FR ($\chi^2(1, N = 427) = 7.44, p < .05$)). Moses, on the other hand, translates impersonal pro-drop better than personal pro-drop (IT→FR ($\chi^2(1, N = 453) = 27.74, p < .05$); ES→FR ($\chi^2(1, N = 427) = 69.14, p < .05$)).

For ES→FR, Its-2 always mistranslates 3rd person singular pro-drop in the source languages as the 1st person singular in French (5a). For IT→FR it translates the impersonal pronoun “*on*” as the 3rd person singular pronoun “*il*” (5b).

- (5) a. ES En dicha empresa *pro* recibía un salario fijo de un millón de dólares al año.
FR Dans cette société **il** recevait un salaire fixe d’un million de dollars par an.
ITS-2 Dans *dicha entreprise **je** recevais un salaire fixe d’un million de dollars au l’an.
EN *In that company he received a fixed salary of a million dollars a year.*
- b. IT Quest’anno si dovrebbe registrare una spesa di 36 miliardi.
FR Cette année, **on** devrait enregistrer une dépense de 36 milliards.
ITS-2 . Cette année **il** devrait enregistrer une dépense de 36 milliards.
EN *This year we should register an expense of 36 billions.*

As for Moses, we expected the translation quality to be lower, since the application domain (news) is very different from the domain of the system’s training data (Europarl), and as expected, many pronouns are missing in the translation because many finite verbs in the source are unknown to the system. Domain adaptation techniques exist, but require additional in-domain training data (Koehn and Schroeder, 2007).

Generally, sentences in this corpus are more complex than sentences in the Europarl corpus. We think this also influences the translation quality in both languages.

6. General Discussion

Results show that there are fewer pro-drop pronouns in the news corpus compared to the parallel and native corpora

¹⁰<http://www.repubblica.it>.

¹¹<http://www.elmundo.es>.

		Its-2			Moses		
Pair	Pro-drop	Correct	Incorrect	Missing	Correct	Incorrect	Missing
IT→FR	personal	44.09%	9.93%	45.96%	18.01%	3.72%	78.26%
	impersonal	13.01%	22.94%	64.04%	42.46%	10.27%	47.26%
	average	24.06%	18.32%	57.61%	33.77%	7.94%	58.27%
ES→FR	personal	37.01%	11.39%	51.60%	12.10%	12.10%	75.80%
	impersonal	23.97%	11.64%	64.38%	48.63%	5.48%	45.89%
	average	32.55%	11.48%	55.97%	24.59%	9.84%	65.57%

Table 6: Percentages of correct, incorrect and missing translation of zero-pronouns. Results obtained by Its-2 and Moses on the news corpus. Average is calculated on the basis of total pro-drop in corpus.

(IT ($\chi^2(2, N = 9\,451) = 84.98, p < .05$); ES ($\chi^2(2, N = 9\,960) = 562.03, p < .05$)). This difference is surprising given that both the news and native corpora are composed of texts produced by native speakers only. The different pro-drop rates can probably be related to the difference in style. News mostly consist of direct and explicit sentences, while texts in native corpus are transcriptions of speeches given at the plenary sessions of the European Parliament.

Both Its-2 and Moses achieve the same performance when translating all three corpora for IT→FR ($\chi^2(2, N = 2\,795) = 0.33, n.s.$). As for ES→FR, there is a difference depending on the corpus: Moses performs better than Its-2 when translating the native corpus, while Its-2 performs better than Moses when translating the news corpus ($\chi^2(2, N = 3\,246) = 31.73, p < .05$).

As already discussed for the parallel and native corpora, the systems' performance decreases when translating impersonal pro-drop. On the news corpus, performance decreases only for Its-2, but Moses achieves a better performance with impersonal pro-drop for both language pairs (for IT→FR: ($\chi^2(1, N = 262) = 57.53, p < .05$); for ES→FR ($\chi^2(1, N = 244) = 43.85, p < .05$)).

Even though we merge incorrect and missing evaluation categories in our χ^2 calculations, we noticed that missing pronouns were a remarkable part of wrong translations, a part even more substantial than incorrect translations themselves ($\chi^2(2, N = 5\,737) = 34.09, p < .05$). Comparing the number of missing personal pro-drop and missing impersonal pro-drop across all three corpora reveals a difference in the systems' performance according to the language pair. For IT→FR, missing impersonal pro-drop pronouns are more frequent in Its-2 translation than personal missing pro-drop pronouns ($\chi^2(2, N = 1\,281) = 16.15, p < .05$). Moses displays no difference between the two categories ($\chi^2(2, N = 1\,034) = 0.93, n.s.$). For ES→FR, on the other hand, missing impersonal pro-drop pronouns are more frequent in the parallel corpus than in the native and news corpora for both systems. Both systems present more missing personal pro-drop in the native and news corpora than in the parallel corpus (for Its-2: ($\chi^2(2, N = 1\,613) = 77.87, p < .05$); for Moses: ($\chi^2(2, N = 1\,092) = 73.83, p < .05$)).

From the high percentages of missing pronouns, it follows that, even more important than the problem of incorrect translations, both systems have a problem recognizing and then generating pro-drop pronouns.

7. Conclusion

In this article we compared the occurrence of null subject pronouns in Italian and Spanish. We evaluated their translation into French, a non pro-drop language, using three different types of corpora: parallel and native from the Europarl corpus, and a third corpus built using newspaper articles.

Null subjects have proven quantitatively important in all three corpora, though their distribution changes depending on the language and the text genre. Spanish has more pro-drop in the parallel and native corpora, and Italian in the news corpus. As for the text genre, the news corpus has fewer pro-drop pronouns than the other two.

From a MT perspective, translation results are determined by the type of pro-drop and the pair of languages involved. Impersonal pro-drop is harder to translate than personal pro-drop, especially for the IT→FR pair. Besides, the choice of corpus affects the results for the ES→FR pair. In summary, Moses outperforms Its-2 on the native corpus; and Its-2 outperforms Moses on the news corpus. This last result is also due to the lack of in-domain training of Moses. The two systems we tested present high percentages of missing pronouns in the target language, especially for impersonal pro-drop, confirming the idea that it is harder to translate impersonal pro-drop than personal pro-drop.

Acknowledgments

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

8. References

- 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.

- Liliane Haegeman. 1994. *Introduction to Government and Binding Theory*. Blackwell Publishers.
- Iustina Ilisei and Claudiu Mihăilă. 2009. The impact of zero pronominal anaphora on translational language: A study on romanian newspapers. In *Proceedings of the Fourth Workshop on Statistical Machine Translation*, pages 90–94.
- Philipp Koehn and Josh Schroeder. 2007. Experiments in domain adaptation for statistical machine translation. In *Proceedings of the Second Workshop on Statistical Machine Translation, StatMT '07*, pages 224–227, Stroudsburg, PA, USA. 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)*.
- Tony McEnery and Richard Xiao. 2008. Parallel and comparable corpora: What is happening? In Gunilla Anderman and Margaret Rogers, editors, *Incorporating Corpora: The Linguist and the Translator*, chapter 2, pages 18–31. Multilingual Matters.
- Ruslan Mitkov. 2002. *Anaphora Resolution*. Longman.
- Marta Recasens and M. Antònia Martí. 2010. AnCora-CO: Coreferentially Annotated Corpora for Spanish and Catalan. *Languages Resources and Evaluation*, 44(4):315–345.
- Luz Rello and Iustina Ilisei. 2009. A Comparative Study of Spanish Zero Pronoun Distribution. In *Proceedings of the International symposium on Data and Sense Mining Machine Translation and Controlled Languages, and their application to emergencies and safety critical domains*, pages 209–214.
- 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)*. European Language Resources Association (ELRA).
- Lorenza Russo, Sharid Loáiciga, and Asheesh Gulati. 2012. Improving machine translation of null subjects in Italian and Spanish. In *Proceedings of the 13th EACL Conference*.
- 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.