

Title of the paper: Bridging the “*gApp*”: improving neural machine translation systems for multiword expression detection

Authors:

Hidalgo-Ternero, Carlos Manuel

University of Malaga (Spain)

Facultad de Filosofía y Letras, Departamento de Traducción e Interpretación

Telephone: +34 661100658

cmhidalgo@uma.es

ORCID ID: <https://orcid.org/0000-0002-8338-2627>

Corpas Pastor, Gloria

University of Malaga (Spain)

Facultad de Filosofía y Letras, Departamento de Traducción e Interpretación

Telephone: +34 952133409

Fax: +34 952131823

gcorpas@uma.es

ORCID ID: <https://orcid.org/0000-0001-6688-1531>

Bridging the “*gApp*”: improving neural machine translation systems for multiword expression detection

Abstract

The present research introduces the tool *gApp*, a Python-based text preprocessing system for the automatic identification and conversion of discontinuous multiword expressions (MWEs) into their continuous form in order to enhance neural machine translation (NMT). To this end, an experiment with semi-fixed Verb-Noun Constructions (VNC) will be carried out in order to evaluate to what extent *gApp* can optimise the performance of the two main free open-source NMT systems —Google Translate and DeepL— under the challenge of MWE discontinuity in the Spanish into English directionality. In the light of our promising results, concluding remarks on how to further optimise MWE-aware NMT systems will round off the chapter.

Keywords: text preprocessing system, Neural Machine Translation (NMT), Multiword Expression (MWE), Verb-Noun Constructions (VNC), discontinuity

1. Introduction

Twenty years after Sag and Baldwin’s (2002) seminal publication, multiword expressions (MWEs) are still *a pain in the neck* for natural language processing systems, and machine translation is not an exception (see the latest survey by Monti et al. 2018). Besides their notoriously problematic features such as syntactic anomaly, non-compositionality, and ambiguity, inter alia, a further challenge arises for Neural Machine Translation (NMT): MWEs do not always consist of contiguous tokens (e.g. *take his piece of advice into consideration*), which seriously hinders their automatic identification and translation (Constant et al. 2017; Foufi et al. 2019; Ramisch and Villavicencio 2018; Rohanian et al. 2019).

In order to overcome the challenges that discontinuous MWEs still pose for even the most robust NMT systems (Zaninello and Birch 2020), we have developed *gApp*, a text preprocessing system for the automatic identification and conversion of discontinuous MWEs into their continuous form, in order to enhance the performance of NMT systems. Against such a background, the overall structure of this paper is as follows. Section 2 describes the literature background of this research. While Section 3 presents the system *gApp*, Section 4 covers the research methodology. In Section 5, the text preprocessing system’s precision and recall will be tested in order to then evaluate to what extent *gApp* can optimise the performance of the two main free open-source NMT systems —Google Translate and DeepL— under the challenge of MWE discontinuity in the Spanish into English directionality. A discussion of the results will then be presented in Section 6. Finally, Section 7 offers concluding remarks on how this tool can enhance the identification and translation of MWEs by NMT systems.

2. Literature background

The quest for optimising the treatment of multiword expressions by current NMT systems has laid a fertile ground for research and debate. In this regard, there is a growing body of literature that has already yielded significant advances in the field of multiword expression detection (Finlayson and Kulkarni 2011; Klyueva et al. 2017; Maldonado et al. 2017; Riedl and Biedman 2016; Zampieri et al. 2019, *inter alia*), and, more precisely, in the automatic identification of discontinuous MWEs (Al Saied et al. 2017, and Al Saied et al. 2019; Alegria et al. 2004; Bejček et al. 2011, and Bejček et al. 2013; Constant et al. 2017; Foufi et al. 2019; Moreau et al. 2018; Rohanian et al. 2019; Schneider et al. 2014), as well as in the optimisation of MWEs' treatment by current NMT systems (Huang et al. 2018; Rikters and Bojar 2017; Wang et al. 2017; Zaninello and Birch 2020).

Against such a background, analogously to *gApp*, there currently exist other MWE processing systems (Alegria et al. 2004; Bejček et al. 2013; Finlayson and Kulkarni 2011; Nagy and Vincze 2014; Ramisch 2015, among other implemented within the framework of the COST action PARSEME, *cfr.* Ramisch et al. 2018) which can also perform a token-based MWE identification by employing a *Lexicon Lookup Method* (Ramisch and Villavicencio 2018), i. e. by resorting to a predefined lexicon of patterns for the automatic detection of MWEs in running text. Among these tools, some of them can specifically identify discontinuous expressions, following analogous strategies to *gApp*: employing a gap-length parameter in order to delimit the maximum number of tokens permitted within MWE constituents (Ramisch 2015), using *surface realisation schemas* (SRS) with the necessary information on the order of the MWE components, their mandatory or optional continuity and their inflectional restrictions (Alegria et al. 2004), and setting the system so that both lexical items in a discontinuous MWE (the preposed and the postposed) must concomitantly trigger the identification process (Foufi et al. 2019), among other techniques.

However, in contrast to all previous state-of-the-art systems, the text preprocessing tool *gApp* brings about a considerable advance since it is not only capable of detecting discontinuous MWEs in running text, but it can also automatically convert them into their continuous form, which significantly improves the performance of, to date, the most robust NMT systems—Google Translate and DeepL—, as it will be demonstrated along the following sections.

3. Overview of the tool

The significant challenges posed by MWE discontinuity have made it necessary to develop *gApp*, a Python-based text preprocessing system for detecting and converting discontinuous MWEs into their continuous form in running text, with a view to enhancing NMT performance. In this regard, *gApp* comprises a lexicon of somatisms (i.e. idioms containing terms that refer to human or animal body parts) conforming semi-fixed Verb-Noun Constructions (VNC) in Spanish. Some examples of somatic VNCs are *sentar la cabeza* (literally, 'to sit the head'; figuratively, 'to settle down') *tomar el pelo* (lit., 'to take someone's hair'; fig., 'to fool someone'), *meter la pata* (lit., 'to put the leg in'; fig., 'to screw it up'), *verse las caras* (lit., 'to see each other's faces'; fig., 'to confront someone'), etc.

For the automatic identification and conversion of MWEs, *gApp* employs the free open-source library *Spacy*, specialised in performing a wide array of advanced NLP tasks, including non-destructive tokenization, POS tagging, dependency parsing, lemmatisation,

and rule-based matching, among others. Regarding rule-based matching, with Spacy it is possible to automatically detect a set of tokens with a specific pattern together with other optional elements that may occur within (and hence split) that MWE (Honnibal and Montani 2017).

Prior to developing identification patterns for *gApp*, a first necessary step was to establish the kind of n-grams that may appear within the discontinuous form of the somatisms under study. To this end, we followed a corpus-based methodology to query two giga-token web-crawled corpora of Spanish (esTenTen and JSI Spanish Timestamped Corpus), both available through Sketch Engine. The esTenTen corpus contains over 17.5 billion words of general Spanish (European and American varieties), while the JSI Spanish Timestamped Corpus comprises over 11.6 billion words of news articles obtained from their RSS feeds (Kilgarriff et al. 2003).

Analogously to Hidalgo-Ternero’s (2020) corpus-based research methodology, Sketch Engine’s Corpus Query Language (CQL) schemas have been applied in order to retrieve both the discontinuous forms of the somatisms under study (henceforth named as *relevant results*) as well as other concordances containing analogous patterns but unrelated to the idiomatic sequences (*irrelevant results*), which will eventually determine the necessary restriction for *gApp*’s detection system so as to optimise its precision and recall. In this regard, the implemented CQL schemas are presented in Table 1.

Table 1. CQL schemas for the discontinuous form of the somatisms

Sequence	CQL schemas
<i>Sentar [1-3 tokens] la cabeza</i>	[lemma="sentar"][] {1,3} [word="la"] [word="cabeza"]
<i>Meter [1-3 tokens] la pata</i>	[lemma="meter"][] {1,3} [word="la"] [word="pata"]
<i>Tomar [1-3 tokens] el pelo</i>	[lemma="tomar"][] {1,3} [word="el"] [word="pelo"]
<i>Verse [1-3 tokens] las caras</i>	[lemma="ver"][] {1,3} [word="las"] [word="caras"]

Once the different concordances with both relevant and irrelevant results were analysed, different rule-based matching patterns were developed in order to identify as many relevant results while filtering out as many irrelevant results as possible. By way of illustration, this has been the base pattern employed to detect discontinuous instances of the somatism *tomar el pelo* with *gApp*:

```
(1) matcher.add('tomar X el pelo', on_match, [{'LEMMA': 'tomar', 'POS': 'VERB'},
      {'POS': {'IN': ['DET', 'ADV', 'ADP', 'PRON', 'PROPN']}}, {'OP': '?'}, {'POS':
      {'NOT_IN': ['VERB', 'CONJ', 'SCONJ', 'ADP']}, 'OP': '?'}, {'ORTH': 'el'}, {'ORTH':
      'pelo'}])
```

Within the method *matcher.add()*, it is possible to include a rule to the matcher, comprising an ID key, a number of patterns and a callback function (in our case, *on_match*) with the arguments *matcher*, *doc*, *id* and *matches*, to act on the matches (Honnibal and Montani 2017). The patterns consist of a list of dictionaries, each of which contains the necessary description of both the exact tokens of the MWE and of those elements that may occur within the sequence. In the case of *tomar el pelo*, the first dictionary of the pattern includes the attribute *LEMMA* for *tomar* with *verb* as its POS tag. For each of the remaining tokens of the MWE (*el pelo*), the attribute *ORTH* had to be added since none of them allow inflection in this idiomatic sequence. Regarding the gap, the concordances retrieved from the analysed corpora reveal that this MWE can be split by

various adverbial, nominal and prepositional phrases such as *más* ('more'), *tanto* ('so much'), *un poco* ('a bit'), *una vez más* ('once again') or any subject, among others. Therefore, the first constituent of the gap was restricted to tokens with the following POS tags in the form of a list: 'DET' ('determiner'), 'ADV' ('adverb'), 'ADP' ('adposition'), 'PRON' ('pronoun'), and 'PROPN' ('proper noun'). For those tokens within the gap whose occurrence is optional, the attribute 'OP': '?' was included. Finally, in order to also filter out those elements at the final position of the gap belonging to the irrelevant results, the code {'POS': {"NOT_IN": ["VERB", "CONJ", "SCONJ", "ADP"]}, 'OP': '?'} had to be incorporated since, according to the concordances retrieved from the analysed corpora, those results containing either a verb, a conjunction, or an adposition directly preceding the bigram *el pelo* were unrelated to discontinuous instances of the somatism *tomar el pelo* and had hence to be excluded so as to increase *gApp*'s precision, as it can be observed in the following instances including irrelevant results:

- (2) Sería la última vez que lo verían como Edwin, pero con marcadas muestras de que Pamela ganaba terreno. "Ya **tomaba hormonas, tenía el pelo** largo y me maquillaba". Luego vendría la operación y el cambio de nombre definitivo, que logró hace unos doce años.
- (3) "Hoy la mujer está súper predispuesta a ver su pelo diferente. Y, sobre todo, a tener un estilo propio", dice Bebe Sanders, estilista de muchas actrices y modelos. "Se está **tomando conciencia de que el pelo** tiene tanto protagonismo en el puntaje estético de una mujer que es casi imperdonable no disfrutarlo", agrega.

Instances 2 and 3 contain concordances with the pattern *tomar* [3 tokens] *el pelo*, which can be considered as irrelevant results because they are literal sequences that are not related to a discontinuous form of the somatism *tomar el pelo*. In instance 2 it is possible to observe the juxtaposition of two main clauses, separated by a comma: "*Ya tomaba hormonas, tenía el pelo largo....*". In this example, both the verb *tomar* and the noun phrase *el pelo* are used in their literal meanings in the two independent clauses, which could hence be translated as "[Pamela] was already taking hormones, had a long hair...". In instance 3, both *tomar* and *el pelo* belong to different clauses joined by the subordinate conjunction *que*: "*Se está tomando conciencia de que el pelo tiene tanto protagonismo...*". In the case of *el pelo* is also employed here with its literal meaning ('hair'), while *tomar* is the verbal collocate of *tomar conciencia* ('to become aware'). This sequence from instance 3 could hence be translated as "*People are becoming aware that hair plays such an important role...*". Despite presenting a similar pattern to the discontinuous form of *tomar el pelo*, these irrelevant results were properly filtered out by *gApp* because in instance 2 the first element of the gap is a noun and the final one is a verb, and in instance 3 the first element of the gap is again a noun and the final one consists of a subordinate conjunction. These elements at the initial and final positions of the gap had been set not to be detected and converted by *gApp* following the predefined patterns for this idiom. Figure 1 provides a partial overview of the lexicon of somatisms and the rule-based matching patterns.

the discontinuous somatisms under study have been retrieved from esTenTen corpus and Spanish JSI Timestamped Corpus, conforming a heterogenous sample in terms of text sources, types and language varieties. In spite of the challenge that noisy input still pose to even the most robust NMT systems (Sperber et al. 2017; Belinkov and Bisk 2018; Anastasopoulos 2019; Niu et al. 2020), concordances containing user-generated content (UGC) were also included in the test so as to mitigate sampling bias, which could otherwise arise from solely analysing NMT canonical training data for the somatisms under study. In this regard, a total of 560 cases was examined, considering the continuous and discontinuous forms of the somatisms *sentar la cabeza*, *meter la pata*, *verse las caras*, and *tomar el pelo*, split by different unigrams, bigrams or trigrams. For each somatism, 70 irrelevant results were compiled, all of which included sequences unrelated to the idiom but whose form may pose some challenges for the automatic identification through *gApp*, in order to calculate, at a first stage, both the precision and recall of this text preprocessing system.

After the calculation of both parameters, at a second stage, the results concerning the NMTs' performance for the different concordances were classified within three main categories depending on whether the somatisms under study were presented in their discontinuous form (category 1) or in their continuous form after being automatically converted through *gApp* (category 2) in contrast to their manual conversion (category 3). DeepL and Google Translate's outputs for these different scenarios were then manually assessed following a reference-based MT evaluation (Hidalgo-Tertero, 2020) with several possible target-text candidates for each of the somatisms in both their continuous and discontinuous forms. In this regard, morphological, syntactic, and/or orthotypographic divergences or source-text/translation inaccuracies affecting other elements in the sentences were not listed *per se* as errors if they were not related to the phenomenon of MWE discontinuity for the somatisms under study.

5. Results

In this section, the results will be analysed and presented at those two stages: firstly, *gApp*'s precision and recall for each of the somatisms under study will be set out in order to assess to what extent *gApp* optimises the performance of the two NMT systems — Google Translate and DeepL— under the challenge of MWE discontinuity.

In the case of *sentar la cabeza*, the tool *gApp* automatically detected and converted 69 forms, 65 of which were true positives and 4 were false positives. In this way, *gApp*'s precision for this idiomatic sequence was 94.2% (65/69 cases) with a recall of 92.9% (65/70 cases). The NMTs' results for the continuous and discontinuous form of this somatism are set out in Figure 2. Regarding DeepL's output, it can be observed that the automatic conversion into the idiom's continuous form through *gApp* did not lead to a significant improvement in this NMT's performance compared to the discontinuous scenario (only 1-case difference, i.e. 1.4%), whereas its manual conversion presented a considerable enhancement (a 6-case difference, i.e. 8.6%). In the case of Google Translate, *gApp* could approach the manual's amelioration with a 5-case divergence, i.e. 7.1%, versus a 7-case one, i.e. 10%, respectively. An illustration of the conversion from the discontinuous to the continuous form of *sentar la cabeza* with *gApp* is presented in Table 2 and the NMTs' performance before and after this conversion through *gApp* is shown in Table 3.

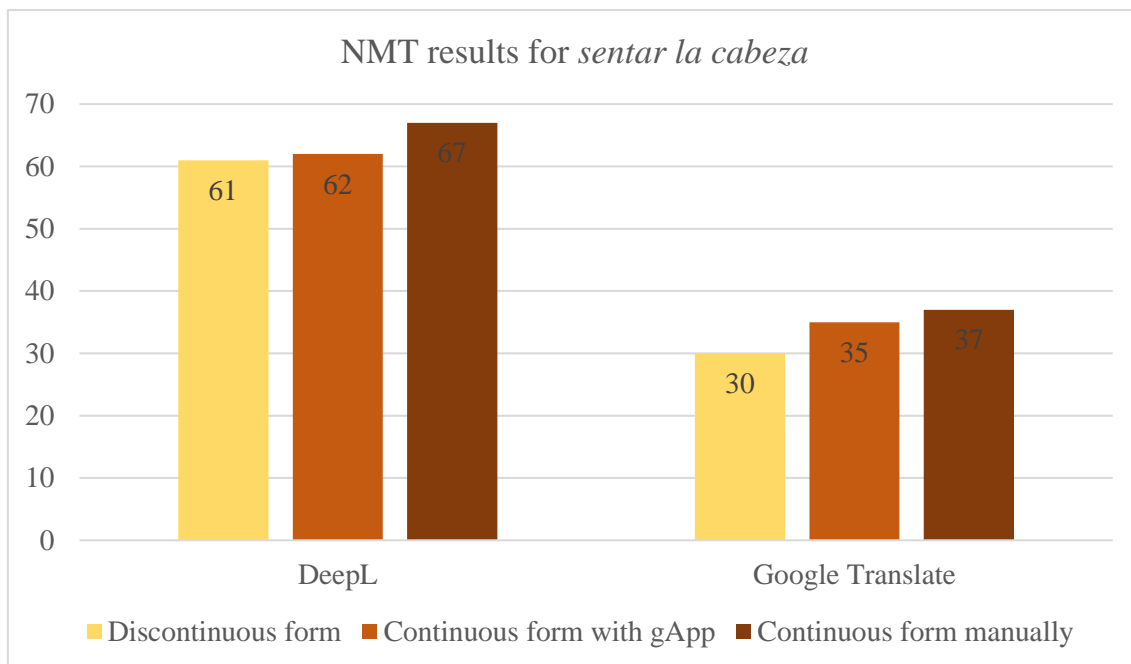


Figure 2. NMT results for *sentar la cabeza*

Table 2. Source-text KWIC extracts with *sentar la cabeza* before and after the conversion with *gApp*

	KWIC extracts
ST [ES] Discontinuous form (original version, before <i>gApp</i>)	Exiliado, Piccolo ha dejado de ser una amenaza para la humanidad y Goku, sin perder su pasión por la lucha, parece haber sentado ligeramente <u>la cabeza</u> tras el nacimiento de su hijo Son Gohan.
ST [ES] Continuous form (after <i>gApp</i>)	Exiliado, Piccolo ha dejado de ser una amenaza para la humanidad y Goku, sin perder su pasión por la lucha, parece haber sentado la cabeza ligeramente tras el nacimiento de su hijo Son Gohan.

Table 3. DeepL and Google Translate’s outcomes before and after the conversion of the ST idiom *sentar la cabeza* with *gApp*.

	NMT outcomes
TT [EN] Discontinuous form in DeepL (before <i>gApp</i>)	In exile, Piccolo is no longer a threat to humanity, and Goku, without losing his passion for fighting, seems to have settled down slightly after the birth of his son Son Gohan.
TT [EN] Continuous form in DeepL (after <i>gApp</i>)	In exile, Piccolo is no longer a threat to humanity, and Goku, without losing his passion for fighting, seems to have settled down slightly after the birth of his son Son Gohan.
TT [EN] Discontinuous form in Google Translate (before <i>gApp</i>)	Exiled, Piccolo is no longer a threat to humanity and Goku, without losing his passion for fighting, seems to have slightly settled his head after the birth of his son Son Gohan.
TT [EN] Continuous form in Google Translate (after <i>gApp</i>)	Exiled, Piccolo is no longer a threat to humanity and Goku, without losing his passion for fighting, seems to have settled down slightly after the birth of his son Son Gohan.

As it can be observed in Table 3, in DeepL both the source-text (ST) continuous and discontinuous idiom was identified and an appropriate equivalent was provided in the target text (TT). However, that was not the case of Google Translate. Before the conversion with *gApp*, this NMT could not properly identify the ST discontinuous idiom, which hence led to a literal and inappropriate translation in the target text. It was only after the conversion with *gApp* that Google Translate could properly detect and offer an adequate equivalent in English for the ST idiom.

As for the somatism *meter la pata*, *gApp* automatically converted 68 forms, all of them corresponding to true positives. Therefore, this tool’s precision amounted to 100% (68/68 cases) and its recall equated to 97.1% (68/70 cases). An overview of the NMTs’ performance for the continuous and discontinuous form of this somatism is presented in Figure 3. When contrasting the data, it is possible to notice that the continuous form of the somatism attained a quasi-analogous outcome for DeepL and Google Translate both through *gApp* (67 and 66 cases, respectively) and the manual conversion (69 and 68). Nevertheless, when comparing it with the discontinuous scenario, major dissimilarities are to be observed: whereas *gApp* and the manual conversion could only result in a slight enhancement for DeepL (2.9% and 5.7%, respectively), both did considerably ameliorate Google Translate’s performance for this idiomatic sequence (12.9% and 15.7% each).

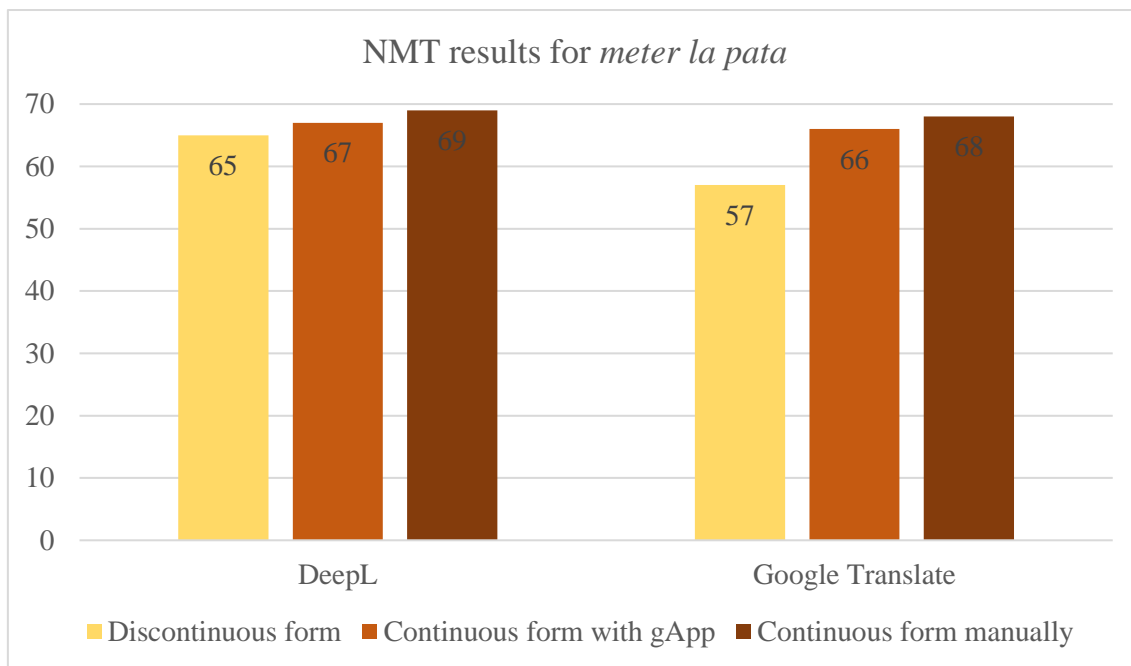


Figure 3. NMT results for *meter la pata*

With regard to *tomar el pelo* (Figure 4), *gApp* transformed 73 cases, including 68 true positives and 5 false positives. In this way, the tool's precision came to 93.2% (68/73) and its recall amounted to 97.1% (68/70 cases) for this somatism. Analogously to the case of *meter la pata*, the continuous form of *tomar el pelo* presented a similar performance for the two NMTs both through the automatic (65 and 62 cases) and the manual (67 and 64) conversion. In contrast to the previous somatism, when comparing the NMTs' performance for the discontinuous and continuous form of *tomar el pelo*, only minor differences have been observed: while, in the case of DeepL, *gApp* ameliorated its performance in a 18.6% and the manual conversion in a 21.4%; Google Translate's outcome was improved in a 15.7% by *gApp* and in a 18.6% by the human conversion. In Table 4 it is possible to observe an instance of the conversion from the discontinuous to the continuous form of *tomar el pelo* with *gApp* and, in Table 5, the NMTs' performance before and after *gApp* is illustrated.

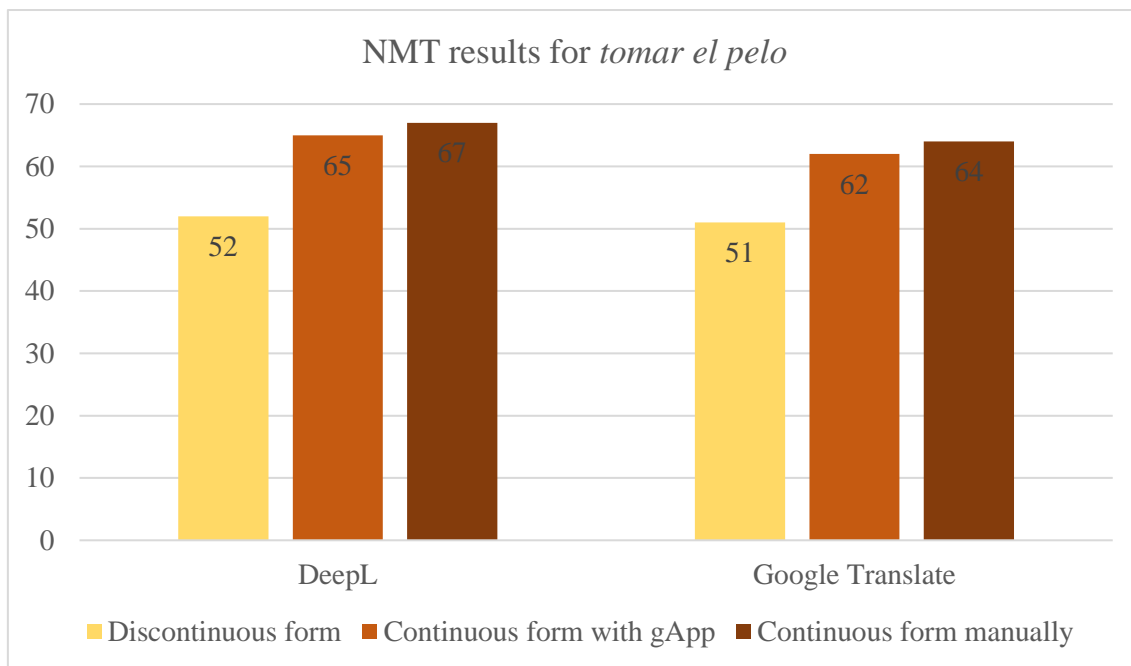


Figure 4. NMT results for *tomar el pelo*

Table 4. Source-text KWIC extracts with *tomar el pelo* before and after the conversion with *gApp*

	KWIC extracts
ST [ES] Discontinuous form (original version, before <i>gApp</i>)	Así consta, ha dicho Monago, en el registro mercantil del Reino Unido, que refleja un reporte anual y el estado financiero de la empresa, según informa el PP extremeño en nota de prensa. </s><s> " ¿Se le puede tomar a la gente el pelo en campaña electoral? ", ha planteado Monago, quien ha señalado que se están riendo de los extremeños diciéndoles que una empresa de 14.000 euros de capital social va a generar tres_millones de turistas al año y va a hacer rascacielos en Castilblanco.
ST [ES] Continuous form (after <i>gApp</i>)	Así consta, ha dicho Monago, en el registro mercantil del Reino Unido, que refleja un reporte anual y el estado financiero de la empresa, según informa el PP extremeño en nota de prensa. </s><s> " ¿Se le puede tomar el pelo a la gente en campaña electoral? ", ha planteado Monago, quien ha señalado que se están riendo de los extremeños diciéndoles que una empresa de 14.000 euros de capital social va a generar tres_millones de turistas al año y va a hacer rascacielos en Castilblanco.

Table 5. DeepL and Google Translate's outcomes before and after the conversion of the ST idiom *tomar el pelo* with *gApp*.

	NMT outcomes
TT [EN] Discontinuous form in DeepL (before <i>gApp</i>)	This is what is stated, said Monago, in the United Kingdom's commercial register, which reflects an annual report and the financial status of the company, according to the Extremadura PP in a press release. </Monago said that they are making fun of the people of Extremadura by telling them that a company with 14,000 euros of share capital will generate three million tourists a year and will build skyscrapers in Castilblanco.
TT [EN] Continuous form in DeepL (after <i>gApp</i>)	This is what is stated, said Monago, in the United Kingdom's commercial register, which reflects an annual report and the financial status of the company, according to the Extremadura PP in a press release. </ Can you make fun of people in an election campaign, " said Monago, who pointed out that they are laughing at the people of Extremadura by telling them that a

	company with 14,000 euros of share capital will generate three million tourists a year and will build skyscrapers in Castilblanco?
TT [EN] Discontinuous form in Google Translate (before <i>gApp</i>)	This is stated, Monago has said, in the commercial register of the United Kingdom, which reflects an annual report and the financial status of the company, as reported by the Extremadura PP in a press release. </s> <s> "Can people tease people in an electoral campaign?" asked Monago, who pointed out that they are laughing at Extremadurans telling them that a company with 14,000 euros of share capital is going to generate three_millions of tourists a year and will make skyscrapers in Castilblanco.
TT [EN] Continuous form in Google Translate (after <i>gApp</i>)	This is stated, Monago has said, in the commercial register of the United Kingdom, which reflects an annual report and the financial status of the company, as reported by the Extremadura PP in a press release. </s> <s> "Can you fool people in the electoral campaign?" asked Monago, who pointed out that they are laughing at Extremadurans telling them that a company with 14,000 euros of share capital is going to generate three_millions of tourists a year and will make skyscrapers in Castilblanco.

The instances in Table 5 show distinctly different results before and after the automatic conversion of the source-text somatism for both NMT systems. In the case of DeepL, the ST discontinuous idiom could not properly be detected, leading to an inappropriate omission of the whole sequence, but after the conversion through *gApp* the continuous idiom could then be identified and an adequate equivalent was provided for the target text. In Google Translate, the ST idiom was detected both in its discontinuous and continuous form. However, before *gApp*, the prepositional phrase *a la gente* ('to the people') following the head verb *tomar* led to its erroneous interpretation both as the subject and the object of the source-text VNC *tomar el pelo*. This prompted an inappropriate translation outcome with the agent and the recipient concomitantly being the same person *"Can people tease people in an electoral campaign?"*. After the conversion through *gApp*, a straightforward analysis of the source text could be performed by Google Translate, which resulted in a more adequate translation outcome *"Can you fool people in the electoral campaign?"*, with an analogous effect to the impersonal ST sentence *"¿Se le puede tomar a la gente el pelo en campaña electoral?"*.

Finally, in the case of *verse las caras* (Figure 5), *gApp* converted 76 forms, 70 of which were true positives and 6 were false positives. Consequently, this tool achieved a precision of 92.1% (70/76 cases) and a recall of 100% (70/70 cases). This recall of all the true-positive cases hence resulted in an equal performance of both the automatic and the manual conversion for DeepL (60 cases each) and Google Translate (48 cases each). The comparison of both the continuous and discontinuous scenarios shows that *gApp* led to a considerable improvement for DeepL (17.1%) and, even more significantly, for Google Translate (22.9%).

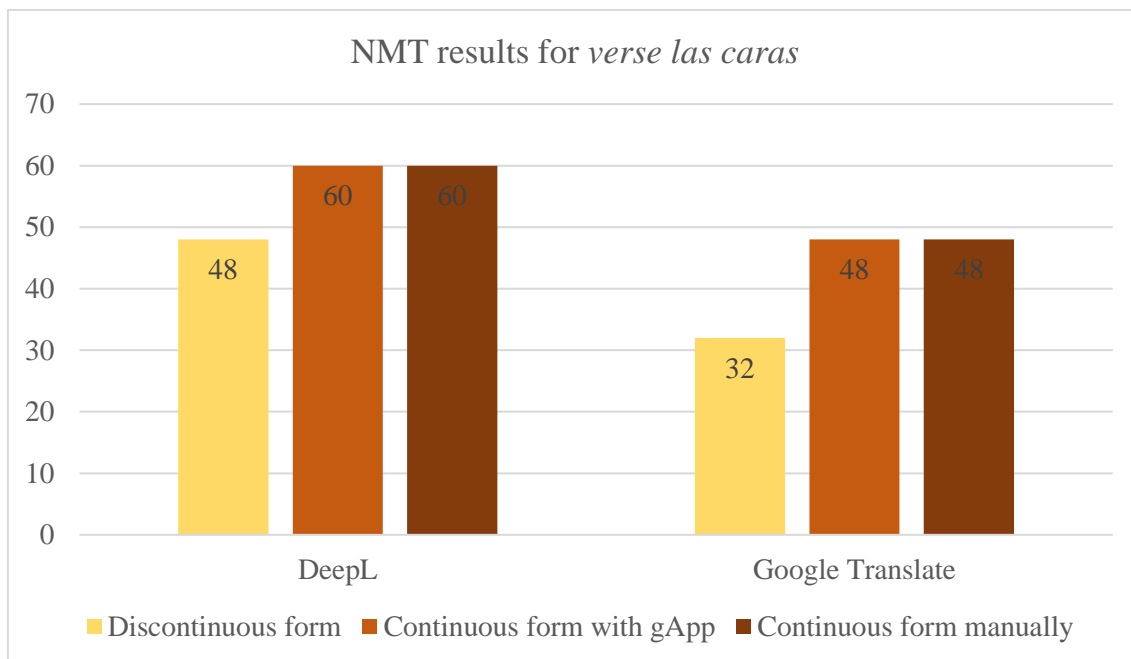


Figure 5. NMT results for *verse las caras*

Once all 560 cases have been analysed and classified, the final results are displayed in Figure 6. For DeepL, the conversion into the continuous form through *gApp* led to an amelioration of 10% (28-case difference), whereas its manual conversion increased its accuracy in a 13.2% (37-case divergence). Regarding Google Translate’s performance, a further differential improvement can be perceived: 14.6% (41 cases) with *gApp* and 16.8% (47 cases) with the human conversion. Finally, the global results (Figure 7), combining both NMTs’ outcomes, show an ultimate contrastive enhancement of 12.3% (69 cases) with *gApp* and of 15% (84 cases) with the manual conversion.

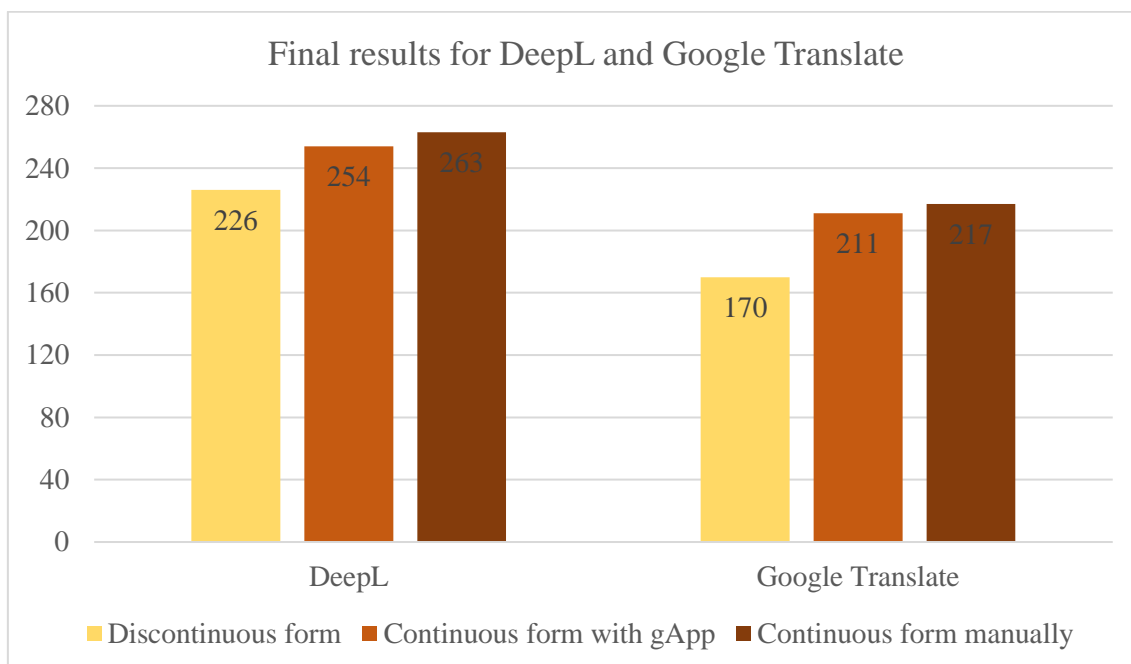


Figure 6. Final results for DeepL and Google Translate

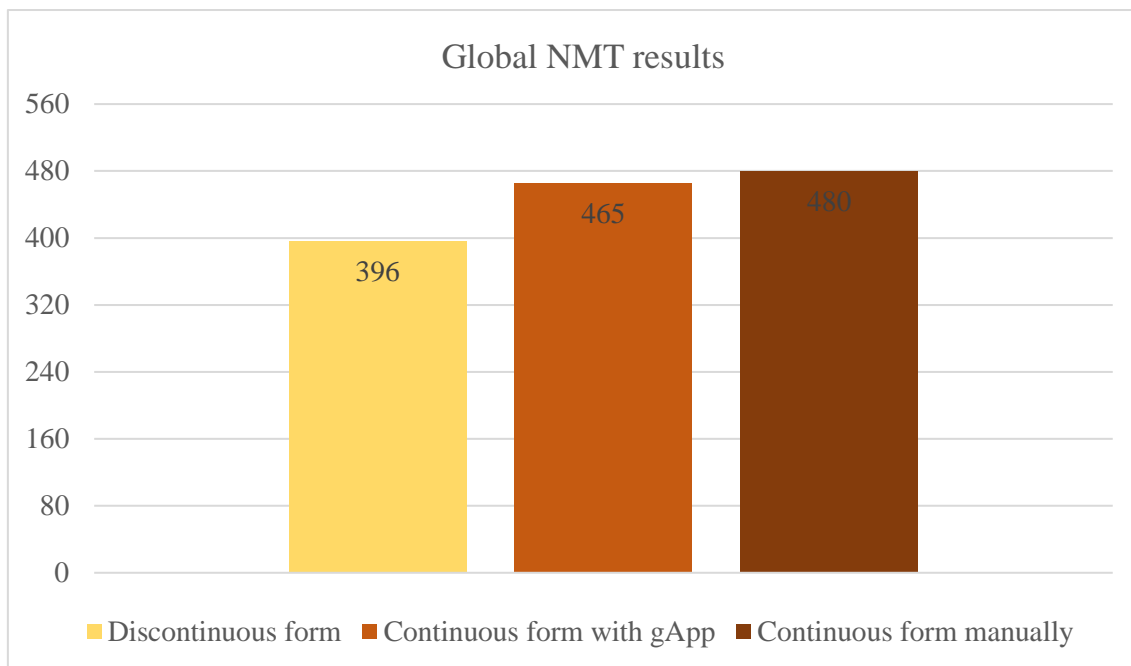


Figure 7. Global NMT results

6. Discussion

The overall NMT results show that *gApp*'s preprocessing system nearly attains the manual conversion for the discontinuous somatisms under study with a 2.7% difference between both of them. This quasi-analogous performance is chiefly due to *gApp*'s refined detection system both in terms of final average precision (94.8%) and recall (96.8%), which means that only 5.2% of the global irrelevant results could penetrate the system and solely 3.2% of all the total relevant results could not be automatically identified.

In order to attain these high accuracy rates in terms of precision and recall some refinements needed to be implemented in *gApp*'s matching system. In spite of its advanced NLP pipeline, Spacy still encountered serious difficulties to automatically detect those cases in which the head verb occurred with an enclitic pronoun, which led to an erroneous POS tagging of these somatisms and, thus, resulted in non-identification. This obstacle hence necessitated the inclusion of regular expressions to also encompass all those relevant results. Despite the insertion of regex, Spacy was still unable to properly detect some instances containing orthotypographic mistakes (such as the absence of written accents when they were due or the non-capitalisation of proper names, inter alia) or rare words (such as *Vd.*, abbreviation of *usted*, i.e. Spanish second person singular pronoun employed in formal forms of address), which emphasises the need for optimising state-of-the-art POS taggers both regarding non-canonical UGC and rare and out-of-vocabulary words (Derczynski et al. 2013; Neunerdt et al. 2013; Gui et al. 2017).

In spite of these limitations in *gApp*'s detection system, the overall NMT results permit us to observe that the tool still offered a considerable enhancement in the translation of somatisms in relative terms, i.e. when comparing the discontinuous versus the continuous scenario after the idiom's automatic conversion, with an average amelioration of 10% (28-case difference) for DeepL and 14.6% (41-case divergence) for Google Translate. Nevertheless, this amelioration was still far from an optimal performance in absolute terms: whereas the continuous form through *gApp* attained a final average accuracy rate of 90.7% (254/280 cases) for DeepL, in the case of Google

Translate its overall performance could only amount to 75.4% (211/280 cases). Two somatisms were particularly challenging for this latter NMT system: *verse las caras*, with an accuracy rate of 68.6% (48/70 cases) for the continuous form through *gApp* and, specially, *sentar la cabeza*, with a final score of 50% (35/70 cases). These results hence accentuate that, despite eliminating the obstacle of discontinuity, a thoroughly accurate performance still necessitates a further optimisation of the NMT systems regarding MWE identification in all its facets.

7. Conclusion

The text preprocessing system *gApp* has proved to optimise NMT performance for the somatisms under study, with an ultimate average enhancement of 12.3% for the analysed NMT systems. It has also proved to attain quasi-analogous results to the human conversion, which in contrast achieved a global 15% improvement. These promising results with semi-fixed VNC somatisms in the Spanish into English directionality invite to further expand *gApp*'s detection lexicon and conversion mechanism in order to examine to what extent it can also lead to NMT enhancement for other MWE categories affected by discontinuity as well as for other NMT systems.

Additionally, the present study can also lay the groundwork for further research to determine the scalation of this model to other language-dependent text preprocessing systems for the automatic conversion of discontinuous MWEs in syntactically flexible languages, in order to enhance MWE-aware NMT systems.

8. Acknowledgements

This paper has been carried out in the framework of various research projects on language technologies applied to translation and interpretation (ref. FFI2016-75831-P, UMA18-FEDERJA-067, CEI-RIS3 and EUIN2017-87746). It has also been funded by the Spanish Ministry of Education (FPU16/02032).

9. References

- Al Saied, Hazem, Mathieu Constant & Marie Candito. 2017. The ATILF-LLF system for Parseme shared task: a transition-based verbal multiword expression tagger. *Proceedings of the 13th Workshop on Multiword Expressions (MWE 2017)*. 127–132.
- Al Saied, Hazem, Marie Candito & Mathieu Constant. 2019. Comparing linear and neural models for competitive MWE identification. *The 22nd Nordic Conference on Computational Linguistics*. 86–96.
- Alegria, Iñaki, Olatz Ansa, Xabier Artola, Nerea Ezeiza, Koldo Gojenola & Ruben Urizar. 2004. Representation and treatment of multiword expressions in Basque. *Proceedings of the Second ACL Workshop on Multiword Expressions: Integrating Processing*. 48–55.
- Anastasopoulos, Antonios. 2019. An analysis of source-side grammatical errors in NMT. *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting*

Neural Networks for NLP. 213–223. <https://www.aclweb.org/anthology/W19-4822> (accessed 15 February 2020)

- Bejček, Eduard, Pavel Stran'ak & Daniel Zeman. 2011. Influence of treebank design on representation of multiword expressions. In Alexander F. Gelbukh (ed.), *Computational Linguistics and intelligent text processing - 12th International Conference, (CICLing 2011)*, vol. 6608 (*Lecture Notes in Computer Science*), 1–14. Berlin, Heidelberg: Springer. DOI:10.1007/978-3-642-19400-9
- Bejček, Eduard, Pavel Stran'ak & Pavel Pecina. 2013. Syntactic identification of occurrences of multiword expressions in text using a lexicon with dependency structures. *Proc. of the 9th Workshop on Multiword Expressions*. 106–115.
- Belinkov, Yonatan & Yonatan Bisk. 2018. Synthetic and natural noise both break neural machine translation. *ArXiv*. <https://arxiv.org/abs/1711.02173> (accessed 2 March 2020)
- Constant, Mathieu, Gülşen Eryiğit, Johanna Monti, Lonneke van der Plas, Carlos Ramisch, Michael Rosner & Amalia Todirascu. 2017. Multiword expression processing: A survey. *Computational Linguistics* 43(4). 1–92.
- Corpas Pastor, Gloria & Jean-Pierre Colson (eds.). 2020. *Computational Phraseology* (IVITRA Research in Linguistics and Literature, 24). Amsterdam & Philadelphia: John Benjamins. ISBN: 978-90-272-0535-3 / E-ISBN: 978-90-272-6139-7.
- Derczynski, Leon, Alan Ritter, Sam Clark & Kalina Bontcheva. 2013. Twitter part-of-speech tagging for all: overcoming sparse and noisy data. *Proceedings of the International Conference on Recent Advances in Natural Language Processing, (RANLP)*. 198–206. <http://www.aclweb.org/anthology/R13-1026> (accessed 5 February 2020)
- Finlayson, Mark & Nidhi Kulkarni. 2011. Detecting multiword expressions improves word sense disambiguation. *Proc. of the ALC Workshop on MWEs (MWE 2011)*. 20–24.
- Foufi, Vasiliki, Luca Nerima & Eric Wehrli. 2019. Multilingual parsing and MWE detection. In Yannick Parmentier & Jakub Waszczuk (eds.), *Representation and parsing of multiword expressions: Current trends*, 217–237. Berlin: Language Science Press.
- Gui, Tao, Qi Zhang, Haoran Huang, Minlong Peng & Xuanjing Huang. 2017. Part-of-speech tagging for twitter with adversarial neural networks. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 2411–2420. <https://www.aclweb.org/anthology/D17-1256.pdf> (accessed 3 March 2020)
- Hidalgo-Ternero, Carlos Manuel. 2020. Google Translate vs. DeepL: analysing neural machine translation performance under the challenge of phraseological variation. In Pedro Mogorrón Huerta & Gloria Corpas Pastor (eds.), *Análisis multidisciplinar del fenómeno de la variación en traducción e interpretación / Multidisciplinary Analysis of the Phenomenon of Phraseological Variation in Translation and Interpreting. MonTI Special Issue 6*.

PARSing and Multiword Expression in Italian), liderado por UNIOR NLP

- Honnibal, Matthew & Inés Montani. 2017, forthcoming. *spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing*.
- Huang, Po-Sen; Chong Wang; Sitao Huang; Denny Zhou & Li Deng. 2018. Towards neural phrase-based machine translation. *International Conference on Learning Representations (ICLR)*. <https://arxiv.org/pdf/1706.05565.pdf> (accessed 17 February 2020)
- Kilgarriff, Adam, Vít Baisa, Jan Bušta, Miloš Jakubíček, Vojtěch Kůvřák, Jan Michelfeit, Pavel Rychlý & Vít Suchomel. 2003. *The Sketch Engine*. <https://www.sketchengine.eu> (accessed 4 March 2020)
- Klyueva, Natalia, Antoine Doucet & Milan Straka. 2017. Neural Networks for Multi-Word Expression Detection. *Proceedings of the 13th Workshop on Multiword Expressions (MWE 2017)*. 60–65. DOI: 10.18653/v1/W17-1707
- Maldonado, Alfredo, Lifeng Han, Erwan Moreau, Ashjan Alsulaimani, Koel Dutta Chowdhury, Carl Vogel & Qun Liu. 2017. Detection of verbal multi-word expressions via conditional random fields with syntactic dependency features and semantic re-ranking. *Proceedings of the 13th Workshop on Multiword Expressions (MWE 2017)*. 114–120. DOI: 10.18653/v1/W17-1715
- Monti, Johanna, Violeta Seretan, Gloria Corpas Pastor & Ruslan Mitkov. 2018. Multiword units in machine translation and technology. In Ruslan Mitkov, Johanna Monti, Gloria Corpas Pastor & Violeta Seretan (eds.), *Multiword Units in Translation and Translation Technology*, 1-37. Amsterdam: John Benjamins.
- Moreau, Erwan, Ashjan Alsulaimani, Alfredo Maldonado & Carl Vogel. 2018. CRF-Seq and CRFDepTree at PARSEME Shared Task 2018: Detecting verbal MWEs using sequential and dependency-based approaches. *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)*. 241–247.
- Nagy T., István & Veronika Vincze. 2014. VPCTagger: Detecting Verb-Particle Constructions with Syntax-Based Methods. *Proceedings of the 10th Workshop on Multiword Expressions (MWE 2014)*. 17–25.
- Neunerdt, Melanie, Bianka Trevisan, Michael Reyer & Rudolf Mathar. 2013. Part-Of-Speech Tagging for Social Media Texts. In Iryna Gurevych, Chris Biemann & Torsten Zesch (eds.), *Language Processing and Knowledge in the Web. Lecture Notes in Computer Science* 8105. 139–150. Berlin, Heidelberg: Springer. DOI: 10.13140/2.1.2905.7284
- Niu, Xing, Prashant Mathur, Georgiana Dinu & Yaser Al-Onaizan. 2020. Evaluating Robustness to Input Perturbations for Neural Machine Translation. *ArXiv*. <https://arxiv.org/abs/2005.00580> (accessed 17 February 2020)

- Ramisch, Carlos. 2015. *Multiword Expressions Acquisition: A Generic and Open Framework. Theory and Applications of Natural Language Processing series XIV*. Switzerland: Springer. ISBN 978-3-319-09206-5.
- Ramisch, Carlos & Aline Villavicencio. 2018. Computational treatment of multiword expressions. In Ruslan Mitkov (ed.), *Oxford Handbook of Computational Linguistics* (2nd ed.). Oxford University Press. DOI: 10.1093/oxfordhb/9780199573691.013.56
- Ramisch, Carlos, Silvio Ricardo Cordeiro, Agata Savary, Veronika Vincze, Verginica Barbu Mititelu, Archana Bhatia, Maja Buljan, Marie Candito, Polona Gantar, Voula Giouli, Tunga Güngör, Abdelati Hawwari, Uxoá Iñurrieta, Jolanta Kovalevskaitė, Simon Krek, Timm Lichte, Chaya Liebeskind, Johanna Monti, Carla Parra Escartín, Behrang QasemiZadeh, Renata Ramisch, Nathan Schneider, Ivelina Stoyanova, Ashwini Vaidya & Abigail Walsh. 2018. Edition 1.1 of the PARSEME Shared Task on Automatic Identification of Verbal Multiword Expressions. *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)*. 222–240. <https://www.aclweb.org/anthology/W18-4925.pdf> (accessed 10 March 2020)
- Riedl, Martin & Chris Biemann. 2016. Impact of MWE resources on multiword recognition. *Proc. of the ACL 2016 Workshop on MWEs*. 107–111.
- Rikters, Matīss & Ondřej Bojar. 2017. Paying attention to multi-word expressions in neural machine translation. *ArXiv*. <https://arxiv.org/abs/1710.06313>
- Rohanian, Omid; Shiva Taslimipour; Samaneh Kouchaki; Le An Ha & Ruslan Mitkov. 2019. Bridging the Gap: Attending to Discontinuity in Identification of Multiword Expressions. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies 1*. 2692-2698.
- Sag, Ivan A., Timothy Baldwin, Francis Bond, Ann Copestake & Dan Flickinger. 2002. Multiword Expressions: A Pain in the Neck for NLP. In Alexander Gelbukh (ed.), *Computational Linguistics and Intelligent Text Processing. CICLing 2002. Lecture Notes in Computer Science*, 1–15. Berlin, Heidelberg: Springer.
- Schneider, Nathan, Emily Danchik, Chris Dyer & Noah A. Smith. 2014. Discriminative lexical semantic segmentation with gaps: Running the MWE gamut. *TACL*, 2. 193–206.
- Sperber, Matthias; Jan Niehues & Alex Waibel. 2017. Toward robust neural machine translation for noisy input sequences. *International Workshop on Spoken Language Translation (IWSLT)*. <https://pdfs.semanticscholar.org/88ed/f12127a628bed608cae0bdf3700d00824df4.pdf> (accessed 13 February 2020)
- Tang, Yaohua, Fandong Meng, Zhengdong Lu, Hang Li & Philip L.H. Yu. 2016. Neural machine translation with external phrase memory. *ArXiv*. <https://arxiv.org/pdf/1606.01792.pdf> (accessed 9 March 2020)

- Wang, Xing, Zhaopeng Tu, Deyi Xiong & Min Zhang. 2017. Translating phrases in neural machine translation. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*. 1421–1431. DOI:10.18653/v1/D17-1149
- Zampieri, Nicolas, Carlos Ramisch & Geraldine Damnati. 2019. The Impact of Word Representations on Sequential Neural MWE Identification. *Joint Workshop on Multiword Expressions and WordNet (MWE-WN 2019)*. 169–175. DOI: 10.18653/v1/W19-5121.
- Zaninello, Andrea & Alexandra Birch. 2020. Multiword expression aware neural machine translation. *Proceedings of the 12th Conference on Language Resources and Evaluation (LREC 2020)*. 3816–3825.