Procesamiento del Lenguaje Natural, Revista nº 51, septiembre de 2013, pp 15-22 recibido 27-04-2013 revisado 16-06-2013 aceptado 21-06-2013

Peeking through the language barrier: the development of a free/open-source gisting system for Basque to English based on apertium.org*

Asomándose por encima de la barrera lingüística: desarollo de un sistema libre/de código abierto euskera–inglés para la asimilación basado en apertium.org

Jim O'Regan
Eolaistriu Technologies
Thurles, Tipperary (Ireland)
joregan@gmail.com

Mikel L. Forcada
Dept. Lleng. i Sist. Informàtics
Universitat d'Alacant
E-03071 Alacant (Spain)
mlf@dlsi.ua.es

Resumen: El artículo describe el desarrollo de un sistema de traducción automática del euskera al inglés pensado para la asimilación (comprensión) construido sobre la plataforma de traducción automática libre/de código fuente abierto basada en reglas Apertium, y lo evalúa preliminarmente usando un nuevo método basado en tests de clausura (*cloze tests*) en los que se pide que se rellenen huecos en una traducción de referencia. Los resultados indican que la disponibilidad de las traducciones en bruto producidas por un sistema con un diccionario de unas 10.000 entradas y unas 300 reglas de traducción incrementan significativamente la capacidad de quien lee para completar los tests con éxito.
Palabras clave: traducción automática, euskera, inglés, software libre, código abierto, asimilación, evaluación

Abstract: The article describes the development of a machine translation system from Basque to English designed for assimilation (*gisting*) built on the free/open-source rule-based machine translation platform Apertium, and evaluates it preliminarly using a new method based in Cloze tests in which readers are asked to fill out gaps in a reference translation. The results indicate that the availability of the raw translations by a system with a dictionary of about 10,000 entries and about 300 translation rules increase significantly the ability of readers to complete the tests successfully.
Keywords: Machine translation, Basque, English, Free Software, Open Source, Gisting, Assimilation, Evaluation

1 Introduction

Machine translation (MT) has become a viable technology that may help individuals in assimilation (to get the gist of a text written in a language the reader does not understand) and dissemination (to produce a draft translation to be post-edited for publication) tasks.

An unfortunate reality for the Basque people on a global scale is that they are little known, and very often for the political conflict and the (fortunately past) violent activity of armed Basque separatist groups. A Basque to English translator would allow the wider world to get a better knowledge about the Basque people and culture through translations of their media (such as the digital edition of the Basque-language daily Berria, http://www.berria.info). Although Google Translate does provide Basque translation, we feel that a free/open-source, rule-based alternative is not only justified, but desirable, for three reasons. Firstly, Google's system often favours fluency over fidelity. Even between larger languages, such as English and Spanish, Google Translate can make errors such as translating *billón* to *million* or removing words from translations (such as *not*) with severe damage to meaning. Secondly, Google's data harvesting has given them something of a 'Big Brother' reputation; to those concerned with privacy, Google

* Development was funded through a grant of the European Association for Machine Translation

ISSN 1135-5948

products are not an option — for a discussion of the risks of online MT see (Porsiel, 2008). Thirdly, having a free/open-source rule-based alternative means that linguistic knowledge (dictionaries, rules) will be explicitly encoded and will therefore be available for reuse in other technologies, thanks to the licence used.

This article describes the development of such a MT system from Basque to English, aimed primarily at assimilation ("gisting"), based on the Apertium[1] free/open-source rule-based MT platform (described in section 2.1), similar to the Basque–Spanish system existing in the project. This involves the creation of a language-pair package which includes monolingual Basque and English dictionaries, a bilingual dictionary, and a set of translation (*structural transfer*) rules. As in the case of the existing Basque to Spanish system (Ginestí-Rosell et al., 2009), the aim was to build a prototype language pair with around 6,000 lemmata in its bilingual dictionary and about 200 structural transfer rules (existing data are described in 2.2; the resulting lexicon and translation rules are described in sections 2.3–2.5. Secondary goals include the creation of the lexical resources necessary for a possible future English to Basque system, and of resources which may be used to augment future work in statistical MT, by using a subset of the rules to generate phrase candidates, to overcome the data sparseness problem (Tyers, 2009).

The evaluation of MT for assimilation purposes is an open research subject which, in our opinion, has not received the attention it deserves in view of the fact that internet-based MT is massively used for assimilation — one could indeed say that assimilation is the most frequent application of MT nowadays. The obvious choice of setting up reading comprehension tests such as multiple choice tests (Jones et al., 2007) is costly and labour-intensive; therefore, alternative ways have been thought. For instance, the evaluation of the Basque to Spanish system (Ginestí-Rosell et al., 2009) involved a two-step procedure very similar to that used in the WMT 2009[2] and WMT 2010[3] workshops on MT (Callison-Burch et al., 2009; Callison-Burch et al., 2010): first, the raw MT output (and nothing else) was shown to target-language monolingual who had to do their best to guess what the sentence meant and post-edit it into a fluent sentence with the same meaning. Then, bilingual people would compare the resulting target sentence with the source sentence and subjectively rate how good a translation it was. This was still a rather costly procedure, required the availability of bilingual experts, and was affected by subjective judgements by bilinguals. For a preliminar evaluation the resulting prototype, we have designed and implemented a novel method for the evaluation of the usefulness of the MT for assimilation or *gisting*, which is described in section 3. To the best of our knowledge, this is the first time that Cloze or gap-filling tests have been used to evaluate the informativeness of MT: note that Cloze tests have indeed been used but in a different way (Somers and Wild, 2000): readability of raw MT was measured by introducing gaps in it and having the subjects fill the gaps. The method described in this paper fills gaps in a reference translation instead. The results indicate that the availability of the raw translations by a system with a dictionary of about 10,000 entries and about 300 translation rules increase significantly the ability of readers to complete the tests successfully.

The paper ends up describing future work (section 4) and giving some concluding remarks (section 5).

2 Development

2.1 The Apertium platform

The system is based on the Apertium MT platform. Originally designed for the Romance languages of Spain, the platform has been extended to support other, more divergent language pairs, including a Basque to Spanish translator (Ginestí-Rosell et al., 2009). Apertium is licensed under the GNU General Public License,[4] as published by the Free Software Foundation, and all software, data, and related source code, for the engine, tools, and all 33 supported language pairs are available to download from the project website.

Apertium uses a shallow-transfer engine. Finite-state transducers are used for lexical processing, hidden Markov models are used for part-of-speech tagging, and finite-state based chunking is used for structural transfer. Linguistic data is encoded in standard-

[1] `http://www.apertium.org`

[2] `http://www.statmt.org/wmt09/`

[3] `http://www.statmt.org/wmt10/`

[4] `http://www.fsf.org/licensing/licenses/gpl.html`

ised, XML-based formats, which are transformed into efficient binary representations using tools which accompany the engine. A full description of the Apertium engine is available in Forcada et al. (2011); what follows is a brief description of the components of the engine.

2.1.1 The pipeline

A typical Apertium-based translation system uses a pipeline of eight components.

- The **deformatter** separates the text to be translated from the formatting markup (e.g., HTML, RTF, OpenDocument, etc.). The format information is encapsulated within "superblanks", which are treated the same as blank characters by the rest of the system.
- The **morphological analyser** tokenises the text into *surface forms*, and adds to each surface form one or more *lexical forms*, containing the *lemma*, and a sequence of tags describing the lexical category and morphological inflection. The analyser is able to analyse both contractions, and fixed-length multi-word units. Multi-word units are processed as "words with spaces", which may themselves be inflected.
- The **part-of-speech tagger** uses a pretrained hidden Markov model (HMM) to determine which of a list of ambiguous analyses is the most likely in the context of its neighbours, delivering a single lexical form.
- The **structural transfer module** uses finite-state pattern matching to detect fixed-length patterns of lexical forms that need special processing due to divergences between the languages: the insertion or deletion of words, reordering, agreement operations (number, gender, person, etc.).
- The **lexical transfer module** is called by the structural transfer module for each source language lexical form, delivering a target-language lexical form.
- The **morphological generator** produces the target-language surface form from each target-language lexical form by inflecting it according to the operations that have been carried out by the structural transfer module.
- The **post-generator** performs orthographical operations, such as contractions.
- Finally, the **reformatter** restores the format information encapsulated by the deformatter module into the translated text, extracting it from the superblanks.

In the case of more complicated language pairs, such as Basque to English, the structural transfer module is split into three components:

- a **chunker**, which performs lexical transfer and local syntactic operations, and segments the input sequence of lexical units into *chunks*; that is, sequences of lexical forms corresponding to some syntactic constituent or feature, such as a verb phrase;
- an **interchunk** module, which performs more global operations on and among chunks; and
- and a **postchunk** module, which performs further local operations on each chunk, before restoring the contents of the chunk to the text stream, to be processed by the morphological generator.

All modules of the system communicate using text streams, via Unix pipes. Individual modules can be replaced, or additional modules added, without any architectural changes to the engine. This allows individual modules to be reused for other natural language processing tasks; for example, the morphological analysis and part-of-speech tagger modules have been reused for opinion classification (Bonev, Ramírez-Sánchez, and Rojas, 2012).

2.2 Existing data

We were able to reuse monolingual data from the Apertium Basque to Spanish system (Ginestí-Rosell et al., 2009), which in turn is based on data from the Matxin (Mayor et al., 2011) Spanish to Basque system.

The Basque morphological dictionary was almost identical to that in the Basque–Spanish system, aside from the addition of extra entries, and the correction of some erroneous entries. Although we added some new inflectional paradigms (for example, for place names ending in *-m* which were

missing from the Basque—Spanish analyser), most of the new paradigms were added for number handling.

The dictionaries for the initial system were built by crossing (or "triangulating") existing Basque to Spanish and Spanish to English lexical data in the Apertium platform using `apertium-dixtools`;[5] the resulting dictionaries were then manually extended for coverage and tuned to yield the most useful lexical equivalents. The creation of the lexicon is further described in section 2.3.

Although the system described here is rule-based, and therefore not corpus-based, it is desirable for many reasons to have a bilingual corpus to aid in the development of the system: to use statistical word alignment tools, to aid in lexicon creation; as a basis for comparison while checking the output of rules; and as a source with which to build a statistical system for the purpose of comparative evaluation; to name but a few. It is also desirable that this corpus be made available under an open-content licence, to be able to distribute it along with the translator[6]

The creation of a bilingual corpus has proved more difficult than expected, as there is hardly any parallel text with a free/open-source licence. Only very recently, two corpora have been made available from OPUS (Tiedemann, 2009),[7] one containing messages from the KDE4 Desktop ("KDE4") and the other one containing film subtitles ("OpenSubtitles2011"); these corpora are not suitable for evaluation: the first one is too dirty and contains only a few complete sentences, and the second one has translations that are very far from being *literal* enough for MT uses. Therefore, two corpora have been built; one is available in the project's repository[8] and contains the Basque and English version of Kafka's *Metamorphosis* - this corpus is cleaner and was used to guide the building of dictionaries, etc., but was not used for evaluation in view of the fact that many translations were found to be problematic; the other corpus was harvested from the website of the International Contact Group for the Basque Country[9] and, after filtering, is a very small corpus (223 sentences); as of today we have not obtained permission to distribute it, but is available on request from the authors.

2.3 Lexicon

The lexicon in the Basque to Spanish translator was based on the lexicon from Matxin, which has a prototype English to Basque translator available from its source repository (Mayor and Tyers, 2009). Despite the availability of the lexicon of this system, we opted not to use it initially: as with Matxin Spanish to Basque, only a small portion of the data is available under a free/open-source licence; of that data, many of the English translations were quite obscure. We chose instead to create our initial lexicon by triangulation, using the Basque to Spanish lexicon and the lexicon from Apertium's English to Spanish translator, and `apertium-dixtools`.

`apertium-dixtools` uses tags, such as those representing lexical category, to restrict the triangulation, thus reducing the number of incorrectly generated entries; that is, the Basque–Spanish entry `izaki<n>:ser<n>` will be aligned with the Spanish–English entry `ser<n>:being<n>`, and not with `ser<v>:be<v>`. In addition, as entries are processed, the software creates a model of the patterns of combinations of tags observed in the source dictionaries, which it sorts by frequency. The crossing model it generates can then be incrementally refined, to generate dictionary entries as close to handwritten entries as possible.

We ran through 6 iterations of this refinement process until we reached the singleton entries which needed to be edited by hand (the majority of these entries were closed category words, though a few were the result of errors in one of the source dictionaries). To keep the Basque–English translator as close as possible to the Basque–Spanish translator, we made use of the greater availability of English–Spanish resources by adding missing entries to the English–Spanish dictionary, and extracted them from a further iteration.

Despite the reduction in triangulation errors achieved by using `apertium-dixtools`, there were still errors due to polysemy and gender ambiguity introduced by the pivot language (for instance, where Spanish nouns in-

[5] `http://wiki.apertium.org/wiki/Apertium-dixtools`

[6] Apertium's development host, Sourceforge, requires that all content distributed via its infrastructure be under a free/open licence.

[7] `http://opus.lingfil.uu.se/`

[8] `http://sourceforge.net/projects/apertium/files/apertium-eu-en/0.3.0/metamorphosis.tmx/download`

[9] `http://hnteuskalherria.org/`

flect like adjectives, such as *hermano* and *hermana*, dictionaries treat the feminine noun as an inflected form of the masculine, leading to an artificial ambiguity).

We used the data from Matxin to automatically resolve *true* polysemy, and to assist in manually resolving the *artificial ambiguity*. Although at the time of development Apertium lacked a lexical selection mechanism for polysemous words,[10] we took this opportunity to annotate the polysemous words we found, for future use. We similarly used Google Translate and online dictionaries to add more data from Matxin (approx. 2,000 words) and from Wikipedia (approx. 1,000 words), though we later discovered that a number of those later additions were duplicates (approx. 500).

The released version of the bilingual lexicon contains about 10,000 entries, of which approx. 350 cannot be reached, because the source-language (Basque) morphological analyser has no corresponding entry. The Basque–Spanish lexicon contains a large number of entries, the majority of which cannot be reached. We estimate it at around 7,000 words.

2.4 Part-of-speech tagger

The Basque part-of-speech tagger in `apertium-eu-es` was used without modification.

2.5 Structural transfer rules

An initial set of structural transfer rules was built by adapting existing rules in the Apertium Basque to Spanish system and then manually extended to deal with the most frequent structural mismatches between English and Basque. The transfer rules were kept as close as possible to the originals, though in most cases the output was quite different, to adapt to English syntax. Most of the new transfer rules were added to handle verb negation, where no rule was required for Spanish. A number of new rules were added to handle phrases expressing dates, and to better handle hyphenated noun–noun constructs.

The existing rules were extended to generate the subject for verbs[11] and to generate auxiliaries; unhyphenated *noun noun* constructs were extended to choose among $noun_1$ $noun_2$, $noun_1$ *'s* $noun_2$ and $noun_2$ *of* $noun_1$ as could be determined. Triple noun phrases were left as-is ($noun_3$ *of* $noun_2$ *of* $noun_1$) as it was the best option, given the data available, though we presume that the availability of further data will show that the possibilities are as varied as for double noun phrases.

[10] A feasible lexical selection module capable of using hand-written rules or rules inferred from a corpus has been only recently described and is in the process of being released (Tyers, Sánchez-Martínez, and Forcada, 2012).

[11] Both Basque and Spanish are pro-drop, so this was not needed in the Basque–Spanish translator.

3 Preliminary evaluation

Since this system was basically aimed at helping English speakers to understand Basque text (assimilation or "gisting"), evaluation has tried to measure the ability of readers to make sense of sentences extracted from documents produced by the International Contact Group cited above. Therefore we recruited (mostly through the mailing list for the Apertium project) 27 people with a good command of English that did not have any command whatsoever of Basque. Of these, 23 people responded by the deadline given.

To that end, we have used a simple method. *Holes* or *gaps* were created in the English version of each sentence by blanking out 20% of the words that were not stopwords to create *Cloze* tests (Taylor, 1953) where subjects had to try their best to guess the missing words and fill them in. In preliminary tests, the 20% percentage was shown to be safely beyond the point where monolingual guessing may be successful. For some randomly chosen sentences, the following hints were given: the original Basque sentence (a weak hint which, however, could be expected to be useful to fill proper names or cognate words), the output of the `apertium-en-es` MT system (which, in case of being maximally useful, would provide information to fill all holes), or both. Each one of the two hints were given with 50% probability separately. Each informer got 32 sentences, roughly 8 in each category.

Instead of using Apertium-eu-en version 0.3.0, released in November, a slightly improved version (subversion repository revision 36906, as of March 24, 2012). Table 2 shows an example where both hints are given.

The sentences rebuilt by the subjects were then compared to the actual English reference sentences and separate success rates were obtained for each of the 4 levels of hinting.

A synonym list (86 entries) was built by manually inspecting the mismatches between words filled in by the informants and those in

Item	Count
Number of bidirectional bilingual dictionary entries	9,565
Number of specific Basque–English bilingual dictionary entries (not to be used in future English–Basque MT)	29
Number of specific English–Basque bilingual dictionary entries (unused in Basque–English MT, ready for future development)	4,727
Number of entries used for Basque–English	9,594
Total dictionary entries	14,321
Structural transfer: chunking rules	197
Structural transfer: inter-chunk rules	55
Structural transfer: post-chunk rules	20
Total structural transfer rules	272

Table 1: Current status of bilingual dictionaries and rules (revision 36906 of `apertium-eu-en`)

Test number 3	
Basque (source language) hint:	`Bruselako Adierazpenaren sinatzaileek argi eta garbi zuzendu dute Adierazpen horrek ordezkatzen duen nazioarteko komunitatearen eskaera.`
Machine translation hint:	`the signatories of the Statement of Brussels clear and clean they have addressed this Statement he of the international community that substitutes the request.`
Problem sentence:	`[sm]@@143: The ##### of the ##### Declaration have addressed in ##### ##### the demands of the international ##### which the ##### Declaration represents.`
Reference sentence:	`The endorsers of the Brussels Declaration have addressed in unequivocal terms the demands of the international community which the Brussels Declaration represents.`

Table 2: Evaluation of `apertium-eu-en`: two types of hints, problem sentence with gaps, and reference sentence with the gaps filled.

the reference sentences. A very conservative approach was taken to select these allowed substitutions (the only substitutions allowed were those not affecting the sense of the reference: actual synonyms in place, alternative spellings, reference-side and informant misspellings, etc.). Table 3 contains some entries of the synonym list.

The number of holes that informants were able to fill out successfully without hints (the baseline) is relatively high. The context provided by other sentences and the homogeneous nature of the corpus (texts related to talks to help the resolution of the conflict in the Basque Country) may be a reason for that. A small increase in success was observed when the Basque source sentence was given as the only hint.

The increase obtained when the output of the Apertium-eu-en machine translation system was presented to the informants is reasonably significant, but lower than expected. Part of the small difference may be due to the fact that the baseline was already rather high.

The most surprising result is that the performance decreases with both the machine-translated and the source sentence are shown to the informants. The decrease is actually quite systematic. This may be due to an "information glut" effect: too much information to integrate.

Hint level	Total number of one-word holes filled	% holes succesfully filled (exactly as in reference)	% holes succesfully filled (using a short synonym list)
No hint	575	26% (sd 13%)	30% (sd 14%)
Source hint (Basque sentence)	543	29% (sd 12%)	34% (sd 14%)
Machine translation hint (Apertium translation)	597	48% (sd 13%)	54% (sd 13%)
Both source and machine translation hint	589	43% (sd 13%)	51% (sd 14%)

Table 4: The results obtained (averages and standard deviations between informants ("sd") are shown).

mesures	measures
mandate	Mandate
likewise	also
legalization	legalisation
lawful	legitimate
laid	set
kept	maintained
international	International
HNT	ICG
financial	economic
evaluation	assessment
end	cessation
encourage	facilitate
demonstrate	prove
declaration	statement
change	transformation
big	major
agents	stakeholders
affected	hit
action	practice
amnisty	amnesty
to	To
truly	sincerely
statemet	declaration
Spain's	Spanish
richest	wealthiest

Table 3: A selection of the synonyms used during the evaluation

When using the synonym list, all results improve slightly, but the trend is the same as when no synonyms are used.

4 Future work

As regards evaluation, this is the first time that MT has been evaluated by measuring its usefulness as a hint in a Cloze test — until now, Cloze tests used for machine translation evaluation asked informants to fill gaps in the MT output itself (Somers and Wild, 2000). However, the experiments reported are still preliminary, and avenues for improvement can easily be identified. For instance, the evaluation corpus was too homogeneous and therefore constituted a very powerful hint, as shown by the results. A general-purpose or multi-domain corpus with wider vocabulary sentences would definitely evaluate it better. Another task that we plan to perform is a comparison task in which Apertium-translated and Google Translate-translated hints were compared.

5 Concluding remarks

This paper has described the building of a free/open-source machine translation system from Basque to English, based on the Apertium platform. The system aims at being a useful tool for assimilation or 'gisting' purposes.

The preliminary results show that it is possible to build, in a few months, a Basque to English MT system capable of producing translations that measurably improve the level of understanding, on the part of non-Basque speakers, of the contents of Basque text.

Acknowledgements: The authors thank the European Association for Machine Translation for financial support, all 23 informants for participating in the evaluation, Francis M. Tyers for useful discussions on the idea of using Cloze tests for machine translation evaluation, and Mireia Ginestí for her help on questions regarding Basque data.

References

Bonev, Boyan, Gema Ramírez-Sánchez, and Sergio Ortiz Rojas. 2012. Opinum: statistical sentiment analysis for opinion classification. In *Proceedings of the 3rd Workshop in Computational Approaches to Subjectivity and Sentiment Analysis*, pages 29–37. Association for Computational Linguistics.

Callison-Burch, Chris, Philipp Koehn, Christof Monz, Kay Peterson, Mark Przybocki, and Omar Zaidan. 2010. Findings of the 2010 joint workshop on statistical machine translation and metrics for machine translation. In *Proceedings of the Joint Fifth Workshop on Statistical Machine Translation and MetricsMATR*, pages 17–53, Uppsala, Sweden, July. Association for Computational Linguistics. Revised August 2010.

Callison-Burch, Chris, Philipp Koehn, Christof Monz, and Josh Schroeder. 2009. Findings of the 2009 Workshop on Statistical Machine Translation. In *Proceedings of the Fourth Workshop on Statistical Machine Translation*, pages 1–28, Athens, Greece, March. Association for Computational Linguistics.

Forcada, Mikel L., Mireia Ginestí-Rosell, Jacob Nordfalk, Jim O'Regan, Sergio Ortiz-Rojas, Juan Antonio Pérez-Ortiz, Felipe Sánchez-Martínez, Gema Ramírez-Sánchez, and Francis M. Tyers. 2011. Apertium: a free/open-source platform for rule-based machine translation. *Machine Translation*, 25(2):127–144.

Ginestí-Rosell, Mireia, Gema Ramírez-Sánchez, Sergio Ortiz-Rojas, Francis M Tyers, and Mikel L Forcada. 2009. Development of a free Basque to Spanish machine translation system. *Procesamiento del Lenguaje Natural*, 43:187–195.

Jones, Douglas, Martha Herzog, Hussny Ibrahim, Arvind Jairam, Wade Shen, Edward Gibson, and Michael Emonts. 2007. ILR-based MT comprehension test with multi-level questions. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 77–80. Association for Computational Linguistics.

Mayor, Aingeru, Iñaki Alegria, Arantza Díaz De Ilarraza, Gorka Labaka, Mikel Lersundi, and Kepa Sarasola. 2011. Matxin, an open-source rule-based machine translation system for Basque. *Machine Translation*, 25(1):53–82.

Mayor, Aingeru and Francis M. Tyers. 2009. Matxin: moving towards language independence. In *Proceedings of the first international workshop on free/open-source rule-based machine translation, Alacant*, pages 11–17.

Porsiel, Jörg. 2008. Machine translation at volkswagen: a case study. *Multilingual Computing & Technology*, 100.

Somers, Harold and Elizabeth Wild. 2000. Evaluating machine translation: the cloze procedure revisited. In *Translating and the Computer 22: proceedings of the Twenty-second International Conference on Translating and the Computer, 16-17 November 2000*.

Taylor, Wilson L. 1953. "Cloze procedure": a new tool for measuring readability. *Journalism Quarterly*, 30:415–433.

Tiedemann, Jörg. 2009. News from opus-a collection of multilingual parallel corpora with tools and interfaces. In *Recent Advances in Natural Language Processing*, volume 5, pages 237–248.

Tyers, Francis M. 2009. Rule-based augmentation of training data in breton-french statistical machine translation. In *Proceedings of the 13th Annual Conference of the European Association of Machine Translation, EAMT09*, pages 213–218.

Tyers, Francis M., Felipe Sánchez-Martínez, and Mikel L. Forcada. 2012. Flexible finite-state lexical selection for rule-based machine translation. In Mauro Cettolo, Marcello Federico, Lucia Specia, and Andy Way, editors, *EAMT 2012: Proceedings of the 16th Annual Conference of the European Association for Machine Translation, Trento, Italy, May 28-30 2012*, pages 213–220.