Automatically Generated Online Dictionaries

Enikő Héja, Dávid Takács

Dept. of Language Technology, Research Institute for Linguistics, HAS
P.O.Box. 360 H-1394, Budapest
eheja@nytud.hu, takdavid@nytud.hu

Abstract

The aim of our software presentation is to demonstrate that corpus-driven bilingual dictionaries generated fully by automatic means are suitable for human use. Need for such dictionaries shows up specifically in the case of lesser used languages where due to the low demand it does not pay off for publishers to invest into the production of dictionaries. Previous experiments have proven that bilingual lexicons can be created by applying word alignment on parallel corpora. Such an approach, especially the corpus-driven nature of it, yields several advantages over more traditional approaches. Most importantly, automatically attained translation probabilities are able to guarantee that the most frequently used translations come first within an entry. However, the proposed technique have to face some difficulties, as well. In particular, the scarce availability of parallel texts for medium density languages imposes limitations on the size of the resulting dictionary. Our objective is to design and implement a dictionary building workflow and a query system that is apt to exploit the additional benefits of the method and overcome the disadvantages of it.

Keywords: Parallel corpus, multi word expressions, dictionary query system

1. Introduction

The objective of this research has been to investigate to what extent LT methods are capable of supporting the creation of bilingual dictionaries. Need for such dictionaries shows up specifically in the case of lesser used languages where it does not pay off for publishers to invest into the production of dictionaries due to the low demand. The targeted size of the dictionaries is between 15,000 and 25,000 entries. Since the completely automatic generation of clean bilingual resources is not possible according to the state of the art, we have decided to provide lexicographers with bilingual resources that can facilitate their work. These kind of lexical resources will be referred to as *proto-dictionaries* henceforward.

After investigating some alternative approaches e.g. hub-and-spoke model (Martin, 2007), alignment of WordNets, we have decided to use word alignment on parallel corpora. Former experiments (Héja, 2010) have proven that word alignment is not only able to help the dictionary creation process itself, but the proposed technique also yields some definite advantages over more traditional approaches. The main motivation behind our choice was that the corpus-driven nature of the method decreases the reliance on human intuition during lexicographic work. Although the careful investigation of large monolingual corpora might have the same effect, being tedious and time-cosuming it is not affordable in the case of lesser used languages.

In spite of the fact that word alignment has been widely used for more than a decade within the NLP community to produce bilingual lexicons e.g. Wu and Xia (1994) and several experts claimed that such resources might also be useful for lexicographic purposes e.g. Bertels et al. (2009), as far as we know, this technique has not been exploited in large-scale lexicographic projects, yet e.g. Atkins and Rundell (2008).

Our earlier experiments has shown that although word alignment has definite advantages over more traditional approaches, there are also some difficulties that have to be dealt with: The method in itself does not handle multi-word expressions and the proto-dictionaries comprise incorrect translation candidates, as well. In fact, in a given parallel corpus the number of incorrect translation candidates strongly depends on the size of the proto-dictionary, as there is a trade-off between precision and recall.

Accordingly, our recent objective is two-fold. First, a proof-of-concept experiment is presented to confirm that the treatment of MWEs is possible within this framework. Secondly, a dictionary query system is described that is apt to exploit the additional benefits of the method and overcome the disadvantages of it. According to our expectations such a system renders the proto-dictionaries helpful for not only lexicographers, but also for ordinary dictionary users. In Section 2. the basic generation process is introduced along with the difficulties we have to deal with. In Section 3. a proof-of-concept experiment is described to confirm that the treatment of MWEs is possible within the proposed framework. The various features of the Dictionary Query System are detailed in Section 4. Finally, a conclusion is given and future work is listed in Section 5.

The proto-dictionaries are available at:

    http://efnilex.efnil.org

2. Generating Proto-Dictionaries – One-Token Translation Pairs

2.1. Input data

Since the amount of available parallel data is crucial for this approach, in the first phase of the project we have experimented with two different language pairs. The Dutch-French language pair represents well-resourced languages while the Hungarian-Lithuaninan language pair represents medium density languages. As for the former, we have exploited the French-Dutch parallel corpus which forms sub-part of the Dutch Parallel Corpus (Macken et al., 2007). It

consists of 3,606,000 French tokens, 3,215,000 Dutch tokens and 186,945 translation units[1] (TUs). As for Hungarian and Lithuanian we have built a parallel corpus comprising 4,189,000 Hungarian and 3,544,000 Lithuanian tokens and 262,423 TUs. Because our original intention is to compile dictionaries covering the every-day language we have decided to focus on literature while collecting the texts. However, due to the scarce availability of parallel texts we made some concessions that might be questionable from a translation point of view. First, we did not confine ourselves purely to the literary domain while collecting texts, that is, the parallel corpus comprises also philosophical works. Secondly, instead of focusing on direct translations between Lithuanian and Hungarian we have relied mainly on translations from a third language. Thirdly, we have treated every parallel text alike, regardless of the direction of the translation, although the DPC contains that information.

2.2. The Generation Process

As already has been mentioned in Section 1., word alignment in itself deals only with one-token units. A detailed description of the generation process of such proto-dictionaries has been given in previous papers, e. g. Héja (2010). In the present paper we confine ourselves to a schematic overview. In the first step the lemmatized versions of each input text have been created by means of morphological analysis and disambiguation[2].

In the second step parallel corpora have been created. We used Hunalign (Varga et al., 2005) for sentence alignment. In the next step word alignment has been performed with GIZA++ (Och and Ney, 2003). During word alignment GIZA++ builds a dictionary-file that stores translation candidates, i.e. source and target language lemmata along with their translation probabilities. We used this dictionary file as the starting point to create the proto-dictionaries.

In the fourth step the proto-dictionaries have been created. Only the most likely translation candidates were kept on the basis of some suitable heuristics, which has been developed while evaluating the results manually. Throughout filtering we have relied on three parameters: the automatically attained *translation probabilities*, and the *corpus frequency data* of the *source* and *target lemmata*.

Finally, the relevant example sentences were provided in a concordance to give hints on the use of the translation candidates.

2.3. Trade-off between Precision and Recall

At this stage of the workflow some suitable heuristics need to be introduced to find the best translation candidates without the loss of too many correct pairs. Therefore, several evaluations were carried out.

---

[1]The size of the parallel corpora is given in terms of translation units instead of in terms of sentence pairs, since many-to-many alignment among source and target sentences was allowed.

[2]The analysis of the Lithuanian texts was performed by the Lithuanian Centre of Computational Linguistics (Zinkevičius et al., 2005). The Hungarian texts were annotated with the tool-chain of the Research Institute for Linguistics, HAS (Oravecz and Dienes, 2002).

It is important to note that throughout the manual evaluation we have focused on lexicographically useful translation candidates instead of perfect translations. The reason behind this is twofold. First, translation synonymy is rare in general language e.g. Atkins and Rundell (2008, p. 467), thus other semantic relations, such as hyponymy or hyperonymy, were also considered. Secondly, since the word alignment method does not handle MWEs in itself, partial matching between SL and TL translation candidates occurs frequently. In either case, provided example sentences make possible to find the right translation.

We considered three parameters when searching for the best translations: The translational probability, the source language lemma frequency and the target language lemma frequency (p_{tr}, F_s and F_t, respectively).

The lemma frequency had to be taken into account for at least two reasons.

1. A minimal amount of data was necessary for the word alignment algorithm to be able to estimate the translational probability.

2. In the case of rarely used TL lemmas the alignment algorithm might assign high translational probabilities to incorrect lemma pairs if the source lemma occurs frequently in the corpus and both members of the lemma pair recurrently show up in aligned units.

Results of the first evaluation showed that translation pairs with relatively low frequency and with a relatively high translational probability yielded cc. 85% lexicographically useful translation pairs. Although the precision was rather convincing, it has also turned out that the size of the resulting proto-dictionaries might be a serious bottleneck of the method (Héja, 2010). Whereas the targeted size of the dictionaries is between 15,000 and 25,000 entries, the proto-dictionaries comprised only 5,521 Hungarian-Lithuanian and 7,007 French-Dutch translation candidates with the predefined parameters. Accordingly, the coverage of the proto-dictionaries should be augmented.

According to our hypothesis in the case of more frequent source lemmata even lower values of translation probability might yield the same result in terms of precision as in the case of lower frequency source lemmata. Hence, different evaluation domains need to be determined as a function of source lemma frequency. That is:

1. The refinement of the parameters yields approximately the same proportion of correct translation candidates as the basic parameter setting,

2. The refinement of the parameters ensures a greater coverage.

Detailed evaluation of the French-Dutch translation candidates confirmed the first part of our hypothesis. We have chosen a parameter setting in accordance with (1) (see Table 1). 6934 French-Dutch translation candidates met the given conditions. 10 % of the relevant pairs was manually evaluated. The results are presented in Table 1. 'OK' denotes the lexicographically useful translation candidates. For instance, the first evaluation range (1st row of Table 1)

comprised translation candidates where the source lemma occurs at least 10 times and at most 20 times in the parallel corpus. With these parameters only those pairs were considered where the translation probability was at least 0.4. As the 1st and 2nd rows of Table 1 show, using different p_{tr} values as cut-off parameters give similar results (87%), if the two source lemma frequencies also differ.

F_s	p_{tr}	OK
$10 \leq LF \leq 20$	p $\geq$ 0.4	83%
$100 \leq LF \leq 200$	p $\geq$ 0.06	87%
$500 \leq LF$	p $\geq$ 0.02	87.5%

Table 1: Evaluation results of the refined French-Dutch proto-dictionary.

The manual evaluation of the Hungarian-Lithuanian translation candidates yielded the same result. We have used this proto-dictionary to confirm the 2nd part of our hypothesis, i.e. that the refinement of these parameters may increase the size of the proto-dictionary. Table 2 presents the results. *Expected* refers to the expected number of correct translation candidates, estimated on the basis of the evaluation sample. 800 translation candidates were evaluated altogether, 200 from each evaluation domain. As Table 2 shows, it is possible to increase the size of the dictionary through refining the parameters: with fine-tuned parameters the estimated number of useful translation candidates was 13,605 instead of 5,521.

F_s	p_{tr}	OK	Expected
$5 \leq LF < 30$	p $>$ 0.3	64%	4,296
$30 \leq LF < 90$	p $>$ 0.1	80%	4,144
$90 \leq LF < 300$	p $>$ 0.07	89%	3,026
$300 \leq LF$	p $>$ 0.04	79%	2,139
			13,605

Table 2: Evaluation results of the refined Hungarian-Lithuanian proto-dictionary.

However, we should keep in mind when searching for the optimal values for these parameters that while we aim at including as many translation candidates as possible, we also expect the generated resource to be as clean as possible. That is, in the case of proto-dictionaries there is a trade-off between precision and recall: the size of the resulting proto-dictionaries can be increased only at the cost of more incorrect translation candidates.

This leads us to the question of what parameter settings are useful for what usage scenarios? We think that the proto-dictionaries generated by this method with various settings match well different user needs. For instance, when the settings are strict so that the minimal frequencies and probabilities are set high, the dictionary will contain less translation pairs, resulting in high precision and relatively low coverage, with only the most frequently used words and their most frequent translations. Such a dictionary is especially useful for a novice language learner. Professional

translators are able to judge whether a translation is correct or not. They might be rather interested in special uses of words, lexicographically useful but not perfect translation candidates, and more subtle cross-language semantic relations, while at the same time, looking at the concordance provided along with the translation pairs, they can easily catch wrong translations which are the side-effect of the method. This kind of work may be supported by a proto-dictionary with increased recall even at the cost of a lower precision. Thus, the Dictionary Query System described in Section 4. in more detail, should be able to support various user needs.

However, user satisfaction has to be evaluated in order to confirm our hypothesis. It forms part of our future tasks.

3. The Treatment of MWEs

3.1. Verbal Structures

A proof-of-concept experiment was performed to confirm that the proposed method might handle MWEs, too[3]. In the first step, the scope of the MWEs was narrowed down to verbal structures. The term verbal structure refers here to expressions consisting of a verb and its salient complementation frame comprising either bound or frequently occurring constituents—regardless if they are complements or adjuncts.

Verbal structures are extracted fully automatically from monolingual corpora. The most important benefit of the method is that it is able to decide automatically when a given nominal head is essentially part of the complementation frame. For instance, in the case of the French expression *donner lieu á* "give rise to" the algorithm detects that both the bound nominal part (*lieu*) and the preposition (*à*) are inherent parts of the verbal structure. In addition, the extraction method seems to be language independent. For a detailed description of the extraction technique see (Sass, 2009).

Table 3 gives some examples of verbal structures, both French and Dutch. In fact, this entry is part of the automatically generated verbal proto-dictionary. Four different Dutch translations were assigned to the French verbal structure *prendre médicament* "take medicine". The most likely Dutch translation is the literary translation *geneesmiddel innemen*. The next translation *geneesmiddel gebruiken* "use medicine" has a slightly different meaning but might be used in similar contexts. Although the third translation *start met gebruik van* "start the use of" is not a whole verbal structure, it can also be considered as a relevant translation. Finally, although the least likely translation *staan onder invloed van drug* – literally: "stand under the influence of drug" – has intuitively a different meaning, one might imagine contexts where the TL expression can be used as the translation of the SL expression.

3.2. Workflow

The workflow is the same as in the case of one-token units except for the fact that a preprocessing model was added

---

[3]For a detailed description of the experiment see: Héja et al. (2011).

$Expression_{source}$	$Expression_{target}$	p_{tr}	F_s	F_t
prendre médicament=obj	neem-in genees-middel=obj	0.377	53	32
	gebruik genees-middel=obj	0.102	53	21
	start gebruik=met:cmp met:cmp-van	0.097	53	28
	sta onder-invloed=particle drug=van:cmp	0.05	53	11

Table 3: A sample entry from the French-Dutch verbal proto-dictionary.

to detect and merge the verbal structures into one-token-expressions so that they could serve as input to GIZA++.

The preprocessing comprised three phases. In the first phase the deep-syntatctic analysis of each side of the Dutch-French parallel corpus was performed. Then the resulting phrase structures were converted to partial dependencies which is the required input format of the next step, where monolingual verbal structures were extracted from each side of the parallel corpus.

Thirdly, the selected verbal structures were merged into one-token expressions in both side of the parallel corpus so that they could serve as input to the alignment algorithm.

The Dutch corpus was parsed by means of Alpino (Bouma et al., 2001) while the French corpus was parsed with the hybrid TIG/TAG parser (Villemonte de la Clergerie, 2010). The extraction technique starts out from the presupposition that the whole verbal structure lies within one clause. Hence, the input corpus has to have clause boundary annotation. Moreover, for a number of reasons the extraction of relevant complementation frames presupposes that noun phrase information is also present in the corpus. On the one hand, we need to know which constituent is the head noun, since it is usually one of the head nouns that is a bound element in the verbal structure, if there is any at all. On the other hand, syntactic markers (prepositions, case suffixes, special syntactic positions) between head nouns and the verb, being part of the complementation frame, should also be detected.

One important decision was made when selecting the features of the deep syntactic annotation that should be taken into account when extracting verbal structures. Two contradictory requirements had to be met: First, the more features are to be considered, the more detailed characterization of the verbal structures is possible. Secondly, exploiting too many features might significanty diminish the results, for irrelevant syntactic labels increase the diversity of the data. Thus, we have experimented with two different feature set.

1. The verb and the head of the direct dependents were kept. Adjectives and complements dependent on the head were also preserved, while determiners were omitted.

2. Only the verbal lemma and the head of the direct object were taken into account. Verbal structures only with different verbal lemma or different direct object were kept separately, otherwise they were merged and the corpus frequencies were recounted.

3.3. Results

The experiment started from the hypothesis that more occurrences of the translated structures in the corpus results in a higher recall with the same precision by setting the translational probability to lower value. Conversely, cutting at the same translational probability, more available data leads to greater precision.

To test this assumption word alignment and dictionary building was run against two different datasets.

The evaluation of the first dataset (whole verbal complementation patterns) resulted in these figures: Cutting at a minimal probability of 0.44 and a minimal source lemma frequency of 100, 46% of the source structures was assigned at least one correct translation and 54% of them was assigned only at least incomplete translations.

In the case of the second dataset with the same parameter setting 60% of the evaluated source lemmata had at least one correct and complete translation. 31% of the source lemmata was assigned only incomplete equivalents, while 9% of the source lemmata had only wrong translations. Our hypothesis proved to be correct only in the case of complete verbal structures, since somewhat surprisingly, the proportion of the correct but incomplete translations dropped significantly in the second case.

4. Dictionary Query System

As earlier has been mentioned, the proposed method has several benefits compared to more traditional approaches:

1. A parallel corpus of appropriate size guarantees that the most relevant translations be included in the dictionary.

2. Based on the translational probabilities it is possible to rank translation candidates ensuring that the most likely used translation variants go first within an entry.

3. All the relevant example sentences from the parallel corpora are easily accessible facilitating the selection of the most appropriate translations from possible translation candidates.

Accordingly, the Dictionary Query System presents some novel features. On the one hand, users can select the best proto-dictionary for their purposes on the Cut Board Page. On the other hand, the innovative representation of the generated bilingual information helps to find the best translation for a specific user in the Dictionary Browser Window.

4.1. Customizable proto-dictionaries: the Cut Board Page

The dictionary can be customized on the Cut Board Page. Two different charts are displayed here showing the distribution of all word pairs of the selected proto-dictionary.

Figure 1: The customized dictionary: the distribution of the Lithuanian-Hungarian translation candidates. Logarithmic frequency of the source words on the x-axis, translation probability on the y-axis.


Figure 2: The customized dictionary: the distribution of the candidates. Logarithmic frequency ratio of the source and target words on the x-axis, translation probability on the y-axis.

1. Plot 1 visualizes the distribution of the logarithmic frequency of the source words and the relevant translation probability for each word pair, selected by the given custom criteria.

2. Plot 2 visualizes the distribution of the logarithmic frequency ratio of the target and source words and the corresponding translation probability for each word pair, selected by the given custom criteria.

Proto-dictionaries are customizable by the following criteria:

1. Maximum and minimum ratio of the relative frequencies of the source and target words (left and right boundary on Plot 1).

2. Overall minimum frequency of either the source and the target words (left boundary on Plot 2).

3. Overall minimum translation probability (bottom boundary on both plots).

4. Several more cut off intervals can be defined in the space represented by Plot 2: Word pairs falling in rectangles given by their left, right and top boundaries are cut off.

After submitting the given parameters the charts are refreshed giving a feedback to the user and the parameters are stored for the session, i. e. the dictionary page shows only word pairs fitting the selected criteria.

4.2. Dictionary Browser

As Figure 3 illustrates, the Dictionary Browser displays four different types of information.

1. List of the translation candidates ranked by their translation probabilities. This guarantees that most often used translations come first in the list (from top to bottom). Absolute corpus frequencies are also displayed.

2. A plot displaying the distribution of the possible translations of the source word according to translation probability and the ratio of corpus ferquency between the source word and the corresponding translation candidate.

3. Word cloud reflecting semantic relations between source and target lemmata. Words in the word cloud vary in two ways.

   First, their *size* depends on their translation probabilities: the higher the probability of the target word, the bigger the font size is.

   Secondly, *colours* are assigned to target words according to their frequency ratios relative to the source word: less frequent target words are cool-coloured (dark blue and light blue) while more frequent target words are warm-coloured (red, orange). Target words with a frequency close to that of the source word get gray colour.

4. Provided example sentences with the source and target words highlighted, displayed by clicking one of the translation candidates.

According to our hypothesis the frequency ratios provide the user with hints about the semantic relations between source and target words which might be particularly important when creating texts in a foreign language. For instance, the Lithuanian lemma *karieta* has four Hungarian eqivalents: "kocsi" (word with general meaning, e.g. 'car', 'railway wagon', 'horse-drown vehicle'), "hintó" ('carriage'), "konflis" ('a horse-drawn vehicle for public hire'), "jármű" ('vehicle'). The various colours of the candidates indicate different semantic relations: the red colour of "kocsi" marks that the meaning of the target word is more general than that of the source word. Conversely, the dark blue colour of "konflis" shows that the meaning of the target word is more special. However, this hypothesis should be tested in the future which makes part of our future work.
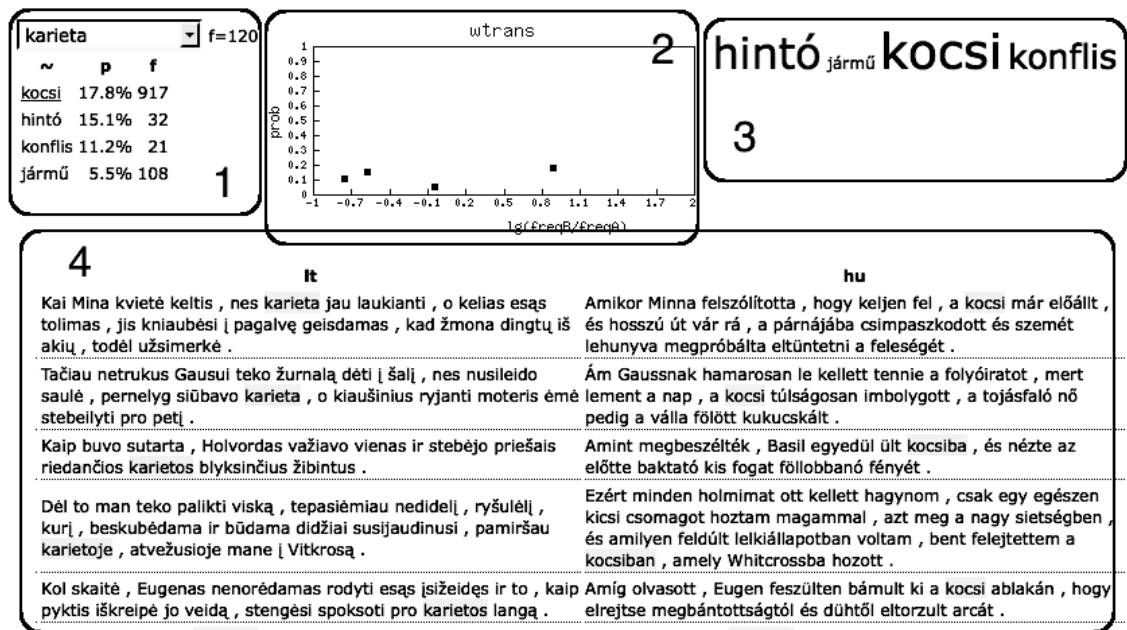
Figure 3: The Dictionary Browser

Figure 4: Querying bigrams

Verbal structures of the form *verb + object* were also uploaded into the query system. As a result, both French and Dutch verbal structures can be queried. As Figure 4 shows, the French word *fiche* "sheet" and "plug" has two different Dutch translations: *stekker* "pin contact, male connector" and *fiche* "sheet". Moreover, the dictionary browser also displays the verbs of which the French noun *fiche* is a frequent object complement. The relevant translations and the concordance can be displayed by clicking any of them.

4.3. Implementation

The online research tool is based on the LAMP web architecture. We use a relational database to store all the data: the multilingual corpus text, sentences and their translations, the word forms and lemmata and all the relations between them. The implementation of such a data structure and the formulation of the queries is straightforward and efficient. The data displayed in the dictionary browser as well as the distributional dataset presented on the charts is selected on-the-fly. The size of the database is linear with the size of the corpus.

5. Conclusions and Future Work

Previous experiments have proven that corpus-driven bilingual resources generated fully by automatic means are apt to facilitate lexicographic work when compiling bilingual dictionaries. In our work this method has been extended to verbal structures and evaluated in two different experimental setup. According to our first results this thread of research looks promising. Given the great variety of possible multiword expressions which can be lexicographically interesting, a more sophisticated approach has to be developed.

We think that the proto-dictionaries generated by this technique with various settings match well different user needs, and consequently, beside lexicographers, they might also be useful for end users, both for language learners and for professional translators. A possible future work is to further evaluate the dictionaries in real world use cases.

Some new assumptions can be formulated which connect the statistical properties of the translation pairs, e.g. their frequency ratios and the cross-language semantic relations between them. Based on the generated dictionaries such hypotheses may be further examined in the future.

In order to demonstrate the generated proto-dictionaries, we have designed and implemented an online dictionary query system, which exploits the advantages of the data-driven nature of the applied technique. It provides different visualizations of the possible translations based on their translation probabilities and frequencies, along with their relevant contexts in the corpus. By presetting different selection criteria the contents of the dictionaries are customizable to suit various usage scenarios.

The dictionaries are publicly available at:

```
http://efnilex.efnil.org.
```

6. Acknowledgements

7. References

Beryl T. Sue. Atkins and Michael Rundell. 2008. *The Oxford Guide to Practical Lexicography*. OUP Oxford.

Ann Bertels, Cédrick Fairon, Jörg Tiedemann, and Serge Verlinde. 2009. Corpus parallèles et corpus ciblés au secours du dictionnaire de traduction. In *Cahiers de lexicologie*, number 94 in Revues, pages 199–219. Classiques Garnier.

Gosse Bouma, Gertjan van Noord, and Robert Malouf. 2001. Alpino: wide-coverage computational analysis of dutch. *Computational Linguistics in the Netherlands 2000. Selected Papers from the 11th CLIN Meeting.*, pages 45–59.

Enikő Héja, Dávid Takács, and Bálint Sass. 2011. Igei bővítménykeretek fordítási ekvivalenseinek kinyerése párhuzamos korpuszból (extracting verbal structures and their translation candidates from parallel corpus). In Attila Tanács and Veronika Vincze, editors, *VIII. Magyar Számítógépes Nyelvészeti Konferencia Tanulmánykötete (Poceedings of the 8th Hungarian Conference on Computational Linguistics)*, pages 47–59, Szeged, Hungary, december.

Enikő Héja. 2010. The role of parallel corpora in bilingual lexicography. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, Mike Rosner, and Daniel Tapias, editors, *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta, may. European Language Resources Association (ELRA).

Lieve Macken, Julia Trushkina, Hans Paulussen, Lidia Rura, Piet Desmet, and Willy Vandeweghe. 2007. Dutch parallel corpus : a multilingual annotated corpus. In *Proceedings of Corpus Linguistics 2007*.

Willy Martin. 2007. Government policy and the planning and production of bilingual dictionaries : The dutch approach as a case in point. *International Journal of Lexicography*, 20(3):221–237.

Franz Josef Och and Hermann Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1):19–51.

Csaba Oravecz and Péter Dienes. 2002. Efficient stochastic part-of-speech tagging for hungarian. In *Proceedings of the Third International Conference on Language Resources and Evaluation*, pages 710–717, Las Palmas.

Bálint Sass. 2009. A Unified Method for Extracting Simple and Multiword Verbs with Valence Information. In *Recent Advances in Natural Language Processing (RANLP 2009)*, pages 399–403, Borovec, Bulgaria.

Dániel Varga, László Németh, Péter Halácsy, András Kornai, Viktor Trón, and Viktor Nagy. 2005. Parallel corpora for medium density languages. In *Recent Advances in Natural Language Processing (RANLP 2005)*, pages 590–596.

Eric Villemonte de la Clergerie. 2010. Convertir des dérivations TAG en dépendances. In *17e Conférence sur le Traitement Automatique des Langues Naturelles - TALN*.

Dekai Wu and Xuanyin Xia. 1994. Learning an english-chinese lexicon from a parallel corpus. In *In Proceedings of the First Conference of the Association for Machine Translation in the Americas*, pages 206–213.

Vytautas Zinkevičius, Vidas Daudaravičius, and Erika Rimkutė. 2005. The Morphologically annotated Lithuanian Corpus. In *Proceedings of The Second Baltic Conference on Human Language Technologies*, pages 365–370.