

# Word Sense Disambiguation as a Wordnets' Validation Method in Balkanet

Dan Tufis<sup>1,2</sup>, Radu Ion<sup>1</sup>, Nancy Ide<sup>3</sup>

<sup>1</sup>Research Institute for Artificial Intelligence Calea „13 Septembrie”, no. 13, PO 050711, Bucharest

<sup>2</sup>University „A.I. Cuza” of Iasi,

<sup>3</sup>Department of Computer Science, Vassar College, Poughkeepsie, NY 12604-0520 USA

<sup>1</sup>tufis@racai.ro, <sup>2</sup>radu@racai.ro, <sup>3</sup>ide@cs.vassar.edu

## Abstract

BalkaNet is a European project which aims at the development of monolingual wordnets for five languages in the Balkans area (Bulgarian, Greek, Romanian, Serbian, and Turkish) and at improvement of the Czech wordnet developed in the EuroWordNet project. The wordnets are aligned to the Princeton Wordnet, according to the principles established by the EuroWordNet consortium. One of the main concerns of this project is the interlingual validation of the wordnets alignment. To this end, we have developed a WSD system based on parallel corpora which exploits the common intuition according to which words that are reciprocal translations in a parallel texts should have the same (or closely related) interlingual meanings. With wordnets under construction our WSD system is mainly a validation tool, pinpointing wrong interlingual alignments, incomplete or missing synsets in one or another of the wordnets.

## Introduction

In previous papers (Ide et al., 2001, 2002) we reported on our sense clustering work, which is based on translation equivalents extracted from parallel corpora. Tufis and Ion (2003) build on this work and further describe a method to accomplish a “neutral” labeling for the sense clusters in Romanian and English that is not bound to any particular sense inventory. Our experiments confirm that the accuracy of word sense clustering based on translation equivalents is heavily dependent on the number and diversity of the languages in the parallel corpus and the language register of the parallel text. For example, using six source languages from three language families (Romance, Slavic and Finno-Ugric), sense clustering of English words was approximately 75% accurate; when fewer languages and/or languages from less diverse families are used, accuracy drops dramatically. This drop is obviously a result of the decreased chances that two or more senses of an ambiguous word in one language will be lexicalized differently in another when fewer languages, and languages that are more closely related, are considered.

To enhance our results, we have explored the use of additional resources, in particular, the aligned wordnets in BalkaNet. BalkaNet is a European project that is developing monolingual wordnets for five Balkan languages (Bulgarian, Greek, Romanian, Serbian, and Turkish) and improving the Czech wordnet developed in the EuroWordNet project. The wordnets are aligned to the Princeton Wordnet, following the principles established by the EuroWordNet consortium. The underlying hypothesis in this experiment exploits the common intuition that reciprocal translations in parallel texts should have the same (or closely related) interlingual meanings (in terms of BalkaNet, ILI record-projections or

simply ILI codes). However, this hypothesis is reasonable if the monolingual wordnets are reliable and correctly linked to the interlingual index (ILI). Quality assurance of the wordnets is a primary concern in the BalkaNet project, and to this end, the consortium developed several methods and tools for validation, described in various papers authored by BalkaNet consortium members (see Proceedings of the Global WordNet Conference, Brno, 2004).

We previously implemented a language-independent disambiguation program, called WSDtool, which has been extended to serve as a multilingual wordnet checker and specialized editor for error-correction. In (Tufis, *et al.*, 2004) it was demonstrated that the tool detected several interlingual alignment errors that had escaped human analysis. In this paper, we describe a disambiguation experiment that exploits the ILI information in the corrected wordnets.

## The Basic Methodology

Our methodology consists of the following basic steps:

1. given a bitext  $T_{L_1L_2}$  in languages  $L_1$  and  $L_2$  for which there are aligned wordnets, extract all pairs of lexical items that are reciprocal translations:  $\{ \langle W_{L_1}^i, W_{L_2}^j \rangle \}$
2. for each lexical alignment  $\langle W_{L_1}^i, W_{L_2}^j \rangle$ , extract the ILI codes for the synsets that contain  $W_{L_1}^i$  and  $W_{L_2}^j$  respectively to yield two lists of ILI codes,  $L_{ILI_1}^1(W_{L_1}^i)$  and  $L_{ILI_2}^2(W_{L_2}^j)$
3. identify one ILI code common to the intersection  $L_{ILI_1}^1(W_{L_1}^i) \cap L_{ILI_2}^2(W_{L_2}^j)$  or a pair of ILI codes  $ILI_1 \in L_{ILI_1}^1(W_{L_1}^i)$  and  $ILI_2 \in L_{ILI_2}^2(W_{L_2}^j)$ , so that  $ILI_1$  and  $ILI_2$  are the *most similar* ILI codes (defined below) among the candidate pairs  $(L_{ILI_1}^1(W_{L_1}^i) \otimes L_{ILI_2}^2(W_{L_2}^j))$  [ $\otimes$  = Cartesian product]

The accuracy of step 1 is essential for the success of the validation method. For this step, we rely on an alignment system that turned in the best performance on English-Romanian in a recent shared task evaluation of word aligners<sup>1</sup> (Tufis, *et al.* 2003) and has since been further improved (Barbu, 2004) to produce the lexicons. The success of step 2 is dependent on the accuracy of the interlingual alignment to find a pair of ILI codes that can disambiguate the translation equivalents.

Our measure of ILI similarity is based on the principle of *hierarchy preservation* (Tufis & Cristea, 2002), which asserts that the *relatedness* (rel) of two ILI records  $R_1$  and  $R_2$  is a measure of *semantic-similarity* (ss) between two synsets  $Syn_1$  and  $Syn_2$  in PWN2.0 that correspond to  $R_1$  and  $R_2$ . We compute semantic-similarity by

$$ss(Syn_1, Syn_2) = 1/1+k$$

where  $k$  is the number of links from  $Syn_1$  to  $Syn_2$  or from both  $Syn_1$  and  $Syn_2$  to the nearest common ancestor. The semantic similarity is 1 when the two synsets are identical (or have the same ILI code), .33 for two sister synsets, and 0.5 for mother/daughter, whole/part, or synsets related by a single link.

Two ILI records  $R_1$  and  $R_2$  are considered closely related if

$$rel(R_1, R_2) = ss(Syn_1, Syn_2) \geq t$$

where  $t$  is an empirical threshold, which in our experiments was set to 0.33 (i.e. we allowed at most two link traversals between what we consider two closely related synsets).

We use a parallel corpus containing texts in  $n+1$  languages (T,  $L_1, L_2 \dots L_k$ ), where for the purposes of disambiguation T is the target language and  $L_1, L_2 \dots L_k$  are the source languages. We also use monolingual wordnets for all  $n+1$  languages, interlinked via an ILI-like structure. The parallel corpus is encoded as a sequence of *translation units* (TU), each containing aligned sentences from each language with tokens tagged and lemmatized as follows<sup>2</sup>:

```
<tu id="Ozz.113">
  <seg lang="en">
    <s id="Oen.1.1.24.2">
      <w lemma="Winston" ana="Np">Winston</w>
      <w lemma="be" ana="Vais3s">was</w>
    ... </s>
  </seg>
  <seg lang="ro">
    <s id="Oro.1.2.23.2">
      <w lemma="Winston" ana="Np">Winston</w>
      <w lemma="fi" ana="Vmii3s">era</w>
    ... </s>
  </seg>
  <seg lang="cs">
    <s id="Ocs.1.1.24.2">
      <w lemma="Winston" ana="Np">Winston</w>
      <w lemma="se" ana="Px---d---ypn--n">si</w>
    ... </s>
  </seg>
  . . .
</tu>
```

<sup>1</sup> [www.cs.unt.edu/~rada/wpt](http://www.cs.unt.edu/~rada/wpt)

<sup>2</sup> For details on the encoding of the corpus, see <http://nl.ijs.si/ME/V2/msd/html/>

For each source language and for all occurrences of a specific word in the target language T, we build a matrix of translation equivalents as shown in Table 1 (eq<sub>ij</sub> represents the translation equivalent in the  $i^{\text{th}}$  source language of the  $j^{\text{th}}$  occurrence of the word in the target language):

	Occ #1	Occ #2	...	Occ #n
$L_1$	eq <sub>11</sub>	eq <sub>12</sub>	...	eq <sub>1n</sub>
$L_2$	eq <sub>21</sub>	eq <sub>22</sub>	...	eq <sub>2n</sub>
...	...	...	...	...
$L_k$	eq <sub>k1</sub>	eq <sub>k2</sub>	...	eq <sub>kn</sub>

**Table 1.** The translation equivalents matrix (EQ matrix)

If the target word is not translated in language  $L_i$ , eq<sub>ij</sub> is represented by the null string. The table is generated as a result of step 1, as described in the previous section.<sup>3</sup> Step 2 transforms the matrix in Table 2 to a matrix with the same dimensions (Table 2) called VSA (Validation and Sense Assignment):

	Occ #1	Occ #2	...	Occ #n
$L_1$	VSA <sub>11</sub>	VSA <sub>12</sub>	...	VSA <sub>1n</sub>
$L_2$	VSA <sub>21</sub>	VSA <sub>22</sub>	...	VSA <sub>2n</sub>
...	...	...	...	...
$L_k$	VSA <sub>k1</sub>	VSA <sub>k2</sub>	...	VSA <sub>kn</sub>

**Table 2.** The VSA matrix

with  $VSA_{ij} = L_{ILI}^{EN}(W_{EN}) \cap L_{ILI}^i(W_{L_i}^j)$ , where  $L_{ILI}^{EN}(W_{EN})$  represent the ILI-codes of all synsets in which the target word  $W_{EN}$  occurs, and  $L_{ILI}^i(W_{L_i}^j)$  is the list of ILI-codes for all synsets in which the translation equivalent for the  $j^{\text{th}}$  occurrence of  $W_{EN}$  occurs.

If no translation equivalent is found in language  $L_i$  for the  $j^{\text{th}}$  occurrence of  $W_{EN}$ ,  $VSA(i,j)$  is undefined; otherwise, it is a set containing 0, 1 or more ILI codes. For undefined VSAs, the algorithm cannot determine the sense number for the corresponding occurrence of the target word. However, it is very unlikely that an entire column in Table 2 is undefined, i.e., that there is no translation equivalent for an occurrence of the target word in any of the source languages.

When  $VSA(i,j)$  contains a single ILI code, the target word occurrence and its translation equivalent are assigned the same sense. For example, the VSA for the English-Romanian translation pair <toe deget> should contain the single ILI-code *ENG20-0528265-n* corresponding to sense 1 of *toe* in PWN and sense number 3 of *deget* in the Romanian wordnet. Thus the disambiguation of this translation pair would be <toe(1) deget(3)>.

When the VSA set is empty—i.e., when none of the senses of the target word corresponds to an ILI code to which a sense of the translation equivalent was linked—the algorithm selects the pair in  $L_{ILI}^{EN}(W_{EN}) \otimes L_{ILI}^i(W_{L_i}^j)$  with the highest ss score. In case of ties, the pair corresponding to the most frequent sense of the target word in the current bitext pair is selected. If this heuristic in turn fails, the choice is made in favor of the pair corresponding to the

<sup>3</sup> In our earlier approach based on clustering, the columns of this table were the vectors used by the agglomerative clustering algorithm (see Ide, *et al.*, 2002).

lowest PWN2.0 sense number for the target word, since PWN senses are ordered by frequency. If no pair in  $L_{ILI}^{EN}(W_{EN}) \otimes L_{ILI}^i(W_{Li}^j)$  meets the semantic similarity requirement, neither the occurrence of the target word nor its translation equivalent can be semantically disambiguated; but once again, it is extremely rare that there is no translation equivalent for an occurrence of the target word in any of the source languages.

When the VSA contains two or more ILI-codes, we have the case of *cross-lingual ambiguity*, i.e., two or more senses are common to the target word and the corresponding translation equivalent in the  $i^{\text{th}}$  language. For example, at least two senses of the English word *movement* are identical to senses of the Romanian word *miscare*. In these cases, the heuristics applied in the case of ties are applied.

### Back-off Method

For the method described in the previous section to succeed, aligned wordnets must be available for all languages in the parallel corpus. Furthermore, it is essential that the quality of the inter-lingual linking is high for all languages concerned. In cases where we cannot fulfill these requirements, we rely on a “back-off” method involving sense clustering based on translation equivalents, as discussed in (Ide, *et al.*, 2002). This method clusters occurrences based on translation equivalents alone, thus eliminating reliance only on high-quality, aligned wordnets. We apply the clustering method after the wordnet-based method has been applied, and therefore each cluster containing an undisambiguated occurrence of the target word will also typically contain several occurrences that have already been assigned a sense. We can therefore assign the most frequent sense assignment in the cluster to previously unlabeled occurrences within the same cluster.

After the method described in the previous section is applied, if a VSA cell contains just one ILI-code, the algorithm successfully sense disambiguates both translation equivalents corresponding to the respective cell (case a above). In all other cases, the algorithm has to make a decision that is occasionally in error. The cases when the algorithm could fail/fails are as follows:

- the VSA cell contains several ILI codes (case b above) or several ILI code pairs (case c above); in this case the algorithm uses the potentially failing tie heuristics presented in the preceding section;
- the VSA cell is empty either because the target word was not translated or because the same word was wrongly translated (by the human translator or by the word-aligner program); in this case, the algorithm as presented in the first section cannot make any decision so, so it fails to disambiguate the senses of the words in the translation equivalence pair.
- the VSA cell is empty because the target word translation does not belong to the source wordnet (it is incomplete); these kind of cells were used to spot the incomplete Romanian synsets when we experimented on the English-Romanian bitext. Following [Kilgarriff, 1997], [Ide et al. 2002], [Tufiş & Ion, 2003] if one does automatically cluster the occurrences of the target word, chances are that probable sense labels would find themselves grouped

with sure ones. The advantages of such a heuristic are:

- it eliminates the requirement of  $k-1$  source wordnets. Indeed, having  $k+1$  languages in our corpus, all we need to apply the WSD method as described before is the aligned wordnets for the target and one source language and alignment lexicons from the target language to every other language in the corpus. The bilingual setting (target language – source language) would ensure the applicability of the WSD procedure and the clustering heuristic would apply a uniform sense labeling among translation equivalents belonging to the same cluster.

it can reinforce or modify the sense assignment for every translation equivalence pair that fall into the cases b) and c) discussed at the end of the previous paragraph, and will be able to assign a sense for all translation pairs falling into the case d), which the previous algorithm was not capable of; all non-disambiguated members of one cluster will be disambiguated according to the majority sense of the already disambiguated members of the cluster;

### Test Data and WSD Evaluation

In order to both evaluate the performance of the WSDtool and assess the accuracy of the interlingual linking of the Balkanet wordnets, we selected a bag of English target nouns, verbs, and adjectives. The words were then extracted from the parallel corpus of George Orwell’s *1984* so that all their senses (at least two per POS) defined in PWN2.0 were also included and interlingually aligned in all Balkanet wordnets, resulting in a set of 211 words which had 1810 occurrences in the English part of the parallel corpus. To create a “gold standard” sense tagging for evaluation purposes, we manually sense-tagged all occurrences of the 211 target words. We then 13 enlisted students enrolled in a Computational Linguistics Masters program to manually assign senses the same occurrences of the target words. An extraction script generated a subset of the 1382 sentences with occurrences of the target words, ensuring that the same sentence was included in at least three student sets. A second set of sense assignments for the 1810 occurrences—each of which was annotated by at least three students—was then established by assigning the sense assigned by the majority of the annotators to each.

Disambiguation results for the same set of words were then generated by the WSDtool algorithm. The system was unable to make a decision for 398 of the 1810 occurrences, primarily in cases where the occurrence was not translated in the Romanian text or was incorrectly determined by the word-aligner.<sup>4</sup> An evaluation program was then applied that generated a file containing detailed information for each of the 1810 occurrences, including

<sup>4</sup> In (Barbu, 2004) we show that the error rate for the latest version of the word aligner (for non-null alignments) is less than 11.5%. This error rate is largely due to English words occurring only once, or English words that are translated differently in each occurrence, so that the corresponding translation pairs are *hapax legomena* that are also not cognates.

- the sense number for that occurrence in the gold standard (GS)
- the majority sense assigned by the student annotators (MAJ)
- the sense assigned by the algorithm(ALG)
- the names of the students who evaluated the occurrence and the sense(s) they assigned

For comparison purposes, we took into account only the 1412 occurrences that were sense disambiguated by the algorithm (without the back-off mechanism). Table 3 summarizes the results. It is interesting to note that the agreement between the algorithm and the gold standard is higher than between the majority and the gold standard.

GS=MAJ	GS=ALG	MAJ=ALG	GS=MAJ=ALG
73.22%	78.68%	67.13%	62.32%

**Table 4.** WSD agreements (without back-off mechanism)

At the present time, the integration of the clustering algorithm with the WSDtool and back-off mechanism evaluation is not completed, and we therefore cannot report results for the fully-implemented method. A rough worst-case estimation of GS=ALG for the full implementation could be conjectured on the basis of the clustering accuracy we reported previously (~75%); therefore, we may assume that the accuracy for the combined WSDTool-cluster method would not be lower than 77%-78%.

## Conclusions

Our disambiguation results using parallel resources are comparable to those obtained in monolingual WSD experiments, but wide-scale application of such a method is, at least at present, impractical, given that parallel resources of sufficient size and covering a substantial variety of languages are scarce. However, more and more parallel resources are becoming available, in particular on the World Wide Web (see for instance <http://www.balkantimes.com> where the same news is published in 10 languages), as well as as a result of the development of WordNets for an increasing number of languages. This opens up the possibility for application of our and similar methods to large amounts of parallel data in the not-too-distant future. One of the greatest advantages of applying such methods to parallel data is that it may be used to automatically sense-tag corpora in not only one language, but rather several at once. The resulting resources could provide substantial training data for monolingual WSD.

## Acknowledgements

The work reported here was carried out within the European project BalkaNet no. IST-2000 29388, with support from the Romanian Ministry of Education and Research under the CORINT programme.

## References

- Barbu A.M. (2004). A word alignment system based on a translation equivalence extractor, this volume.
- Erjavec, T., Lawson A., Romary, L.(eds.) (1998). East Meet West: A Compendium of Multilingual Resources. TELRI-MULTEXT EAST CD-ROM, 1998, ISBN: 3-922641-46-6.

Ide N., Erjavec T., Tufis, D. (2001). Automatic Sense Tagging Using Parallel Corpora. Proceedings of the 6<sup>th</sup> Natural Language Processing Pacific Rim Symposium, Tokyo, 83-9.

Ide, N., Erjavec, T., Tufis, D (2002). Sense Discrimination with Parallel Corpora. Proceedings of the SIGLEX Workshop on Word Sense Disambiguation: Recent Successes and Future Directions. ACL2002, Philadelphia, 54-60.

Tufis, D., Cristea, D. (2002). Methodological issues in building the Romanian Wordnet and consistency checks in Balkanet. Proceedings of LREC2002 Workshop on Wordnet Structures and Standardisation, Las Palmas, May, 35-41

Tufis, D., Barbu, A.M. (2002): Revealing translators knowledge: statistical methods in constructing practical translation lexicons for language and speech processing. International Journal of Speech Technology. Kluwer Academic Publishers, no. 5, 25-30,

Tufis, D. (2002) .A cheap and fast way to build useful translation lexicons. Proceedings of the 19th International Conference on Computational Linguistics, COLING2002, Taipei, 1030-1036.

Tufis D., Ion R (2003). Word sense clustering based on translation equivalence in parallel texts; a case study in Romanian. Proceedings of the International Conference on Speech and Dialog SPED, Bucharest (2003).

Tufis D., Barbu A.M., Ion R. (2003). A word-alignment system with limited language resources. Proceedings of the NAACL 2003 Workshop on Building and Using Parallel Texts; Romanian-English Shared Task, Edmonton, 36-39.

Tufis, D., Ion, R., Barbu, E. Barbu, V. (2004). Cross-Lingual Validation of Wordnets. Proceedings of the 2<sup>nd</sup> International Wordnet Conference, Brno.