

## Chapter 3

# Enriching Slovene wordnet with domain-specific terms

Špela Vintar

Darja Fišer

Dept. of Translation, Faculty of Arts, University of Ljubljana

The paper describes an innovative approach to expanding the domain coverage of the Slovene wordnet (sloWNet) by exploiting multiple resources. In the experiment described here we are using a large monolingual Slovene corpus of texts from the domain of informatics to harvest terminology from, and a parallel English-Slovene corpus and an online dictionary as bilingual resources to facilitate the mapping of terms to sloWNet. We first identify the core terms of the domain in English using the Princeton University's WordNet 2.1, and then we translate them into Slovene using a bilingual lexicon produced from the parallel corpus. In the next step we extract multi-word terms from the Slovene domain-specific corpus using a hybrid approach, and finally match the term candidates to existing wordnet synsets. The proposed method appears to be a successful way to improve the domain coverage of the wordnet as it yields abundant term candidates and exploits various multilingual resources.

## 1 Introduction

WordNet (Fellbaum 1998) is an extensive lexical database in which words are divided by part of speech and organized into a hierarchy of nodes. Each node represents a concept, and words denoting the same concept are grouped into a synset with a unique id (e.g. ENG20-02853224-n: {car, auto, automobile, machine, motorcar}). Concepts are defined by a short gloss (e.g. 4-wheeled motor vehicle, usually propelled by an internal combustion engine) and are also linked to other relevant synsets in the database (e.g. hypernym: {motor vehicle, automotive vehicle}, hyponym: {cab, hack, taxi, taxicab}). Over time, WordNet has become one of the most valuable resources for a wide range of natural language



processing applications, which initiated the development of wordnets for many other languages as well.<sup>1</sup>

One of such enterprises is the building of sloWNet, the Slovene wordnet (Erjavec & Fišer 2006; Fišer 2007; Fišer & Sagot 2008). While this task would normally involve substantial manual labour and the efforts of several linguists, sloWNet was built almost single-handedly exploiting multiple multilingual resources, such as bilingual dictionaries, parallel corpora and online semantic resources.

A combination of all these approaches yielded the first version of the Slovene wordnet<sup>2</sup> (sloWNet) containing about 17,000 synsets and 20,000 literals. However, the majority of these literals are single-word items, because the main lexicon extraction procedures involved in the building of a wordnet involved no systematic handling of multi-word expressions. Also, sloWNet can only be as good as the resources that had been used for its construction. While the coverage for some domains, such as botany or zoology, is excellent, other domains remain underrepresented with numerous lexical gaps still to be filled. If we wish to use a wordnet in any domain-specific application, such as Word Sense Disambiguation or Machine Translation, it is crucial that it contains the terminology of the target domain. The purpose of this paper is to propose a method to enrich a wordnet with domain-specific single- and multi-word expressions.

The target domain in the experiments described below is information technology (IT), for which we have a 15 million word monolingual corpus and a small 300,000 word parallel corpus. We use automatic term recognition to extract multi-word IT terms from the large Slovene corpus and word alignment to extract a bilingual lexicon of single-word terms from the parallel corpus. Using this lexicon and a domain-specific bilingual dictionary as a bridge across the two languages we connect the Slovene multi-word terms to the wordnet hierarchy via English, ie. the Princeton WordNet (PWN).

The rest of the paper is organized as follows: first, the sloWNet Project is described. §3 describes the resources used and the procedure to extract domain-specific expressions from the corpus. §4 presents the bilingual part of the experiment where we try to map terms to the wordnet hierarchy. The results are discussed and evaluated in §5, and the paper ends with concluding thoughts and plans for future work.

---

<sup>1</sup> See [http://www.globalwordnet.org/gwa/wordnet\\_table.htm](http://www.globalwordnet.org/gwa/wordnet_table.htm)

<sup>2</sup> sloWNet is distributed under the Creative Commons licence, <http://nl.ijs.si/slownet>

## 2 Building sloWNet

The first version of the Slovene wordnet was created on the basis of the Serbian wordnet (Krstev et al. 2004), which was translated into Slovene with a Serbian-Slovene dictionary. The main advantages of this approach were the direct mapping of the obtained synsets to wordnets in other languages and the density of the created network. The main disadvantage was the inadequate disambiguation of polysemous words, therefore requiring extensive manual editing of the results. The core sloWNet contains 4,688 synsets, all from Base Concept Sets 1 and 2.

In the process of extending the core sloWNet we tried to leverage the resources we had available, which are mainly corpora. Based on the assumption that translations are a plausible source of semantics we used multilingual parallel corpora such as the Multext-East (Erjavec & Ide 1998) and the JRC-Acquis corpus (Steinberger et al. 2006) to extract semantically relevant information (Fišer 2007).

We assumed that the multilingual alignment based approach can either convey sense distinctions of a polysemous source word or yield synonym sets based on the following criteria (cf. Dyvik 1998, Diab & Resnik 2002 and Ide, Erjavec & Tufis 2002):

(a) senses of ambiguous words in one language are often translated into distinct words in another language (e.g. Slovene equivalent for the English word *school* meaning ‘educational institution’ is *šola* and *jata* for a large group of fish);

(b) if two or more words are translated into the same word in another language, then they often share some element of meaning (e.g. the English word *boy* meaning a ‘young male person’ can be translated into Slovene as either *fant* or *deček*).

In the experiment, corpora for up to five languages (English, Slovene, Czech, Bulgarian and Romanian) were word-aligned, with Uplug (Tiedemann 2003) used to generate a multilingual lexicon that contained all translation variants found in the corpus. The lexicon was then compared to the existing wordnets in other languages. For English, the Princeton WordNet (Fellbaum 1998) was used while for Czech, Romanian and Bulgarian, wordnets developed in the BalkaNet project (Tufis 2000) were used. If a match between the lexicon and wordnets across all the languages was found, the Slovene translation was assigned the appropriate synset id. In the end, all the Slovene words sharing the same synset ids were grouped into a synset.

The results obtained in the experiment were evaluated automatically against a manually created gold standard. A sample of the generated synsets was also checked by hand. The results were encouraging, especially for nouns with f-

measure ranging between 69 and 81%, depending on the datasets and settings used in the experiment. However, the approach had two serious limitations: first, the automatically generated network contains gaps in the hierarchy where no match was found between the lexicon and the existing wordnets, and second, the alignment was limited to single-word literals, thus leaving out all the multi-word expressions.

We tried to overcome this shortcoming with extensive freely available multi-lingual resources, such as Wikipedia and Eurovoc. These resources are rich in specialized terms, most of which are multi-word. Since specialized terminology is typically monosemous, a bilingual approach sufficed to translate monosemous literals from PWN 2.0 into Slovene. A bilingual lexicon was extracted from Wikipedia, Wiktionary and Wikispecies by following inter-lingual links that relate two articles on the same topic in Slovene and English. We improved and extended this lexicon with a simple analysis of article bodies (capitalisation, synonym extraction, preliminary extraction of definitions). In addition we extracted a bilingual lexicon from Eurovoc, a multilingual thesaurus that is used for classification of EU documents. This procedure yielded 12,840 synsets. The translations of the monosemous literals are very accurate and include many multi-word expressions, and thus neatly complement the previous alignment approach. Also, they mostly contain specific, non-core vocabulary.

Synsets obtained from all three approaches were merged and filtered according to the reliability of the sources of translations. The structure of PWN synsets for which no translation could be found with any of the approaches was adopted from PWN based on the hierarchy preservation principle (Tufis 2000), only the literals were left empty. The entire network of synsets was then formatted in DE-BVisDic XML (Horák et al. 2005). The latest version of sloWNet (2.1, 30/09/2009) contains about 20,000 unique literals, which are organized into almost 17,000 synsets, covering about 15% of PWN. Base Concept Sets 1 & 2 are fully covered but there are also many specific synsets. The most frequent domain in sloWNet is Factotum (25%) which was mostly obtained from the dictionary and a parallel corpus, while the following three are Zoology (17%), Botany (13%) and Biology (7%) and come from Wikipedia.

sloWNet mostly contains nominal synsets (91%), and there are some verbal and adjectival synsets as well. Apart from single word literals, there are also quite a few multi-word expressions (43%). These too mostly come from Wikipedia. Synsets in sloWNet are relatively short as 66% of them contain only one literal, average synset length being 1.16 literals. The longest synset contains 16 literals (for verb *goljufati* ‘cheat’). The most common relation in sloWNet is hypernymy,

which represents almost half of all relations in the wordnet (46%). Hypernymy is by far the most prevalent relation for nouns (91%). Nominal hypernymy chains tend to be quite long, the longest ones containing 16 synsets. Since sloWNet does not cover the entire inventory of PWN concepts, there are some gaps (empty synsets) in the network. An investigation of nominal hierarchies revealed that almost half (46%) of the chains do not contain a single gap and that there are only 2% of chains with five or more gaps. These gaps will have to be filled in the future in order to obtain a denser hierarchy of nodes.

### 3 Harvesting domain-specific terminology from specialised corpora

#### 3.1 Multi-word expressions and wordnet

Multi-word expressions (MWE) are lexical units that include a range of linguistic phenomena, such as nominal compounds (e.g. *blood vessel*), phrasal verbs (e.g. *put up*), adverbial and prepositional locutions (e.g. *on purpose*, *in front of*) and other institutionalized phrases (e.g. *de facto*). MWEs constitute a substantial part of the lexicon, since they express ideas and concepts that cannot be compressed into a single word. Moreover, they are frequently used to designate complex or novel concepts. As can be seen in Table 1, the majority of MWEs in the Princeton WordNet do not belong into any of the Basic Concept Sets, meaning that they encode specialized concepts and are frequently terms.

As a consequence, their inclusion into a wordnet is of crucial importance, because any kind of semantic application without appropriate handling of MWEs is severely limited.

Table 1: The distribution of MWEs in PWN across BCS

Group	Frequency
other	64,205
BCS 3	1,470
BCS 2	926
BCS 1	339
total	66,940

For the purpose of MWE identification, various syntactical (Bourigault 1993), statistical (Tomokiyo & Hurst 2003) and hybrid semantic-syntactic-statistical methodologies (Piao et al. 2003, Dias & Nunes 2004) have been proposed, to name but a few. Since the majority of MWEs included in the Princeton WordNet are nominal (see Table 2) and compositional, our approach is based on syntactic features of MWEs.

Table 2: The distribution of MWEs in PWN across part-of-speech (The figures were taken from Princeton WordNet 2.1.)

POS	Frequency
nouns	60,931
verbs	4,315
adverbs	955
adjectives	739
total	66,940

In addressing the issue of MWEs in sloWNet, we initially wanted to find Slovene equivalents for the MWEs already present in Princeton WordNet. We describe this experiment and its successful implementation in (Vintar & Fišer 2008).

### 3.2 Resources

If a wordnet is to be used in a semantic application within a specific domain, we wish to ensure its coverage within this domain primarily for the target language. The goal we address here is thus how to enrich sloWNet with domain-specific Slovene MWEs regardless of whether their English counterparts are included in PWN or not.

The resources we use to this end are the following (Figure 1):

- Ikorpus, a Slovene corpus of Computer Science texts, size ca. 15 million words, morphosyntactically annotated and lemmatized,
- a Slovene-English parallel corpus of Computer Science abstracts, size ca. 300,000 words, morphosyntactically annotated and lemmatized,
- Islovar, a Slovene-English online dictionary of Computer Science,
- Princeton WordNet.

The idea underlying our approach is that a large domain-specific corpus, especially one sufficiently varied in terms of register and text types, can be an excellent source of domain knowledge. Using terminology extraction, gloss extraction and relation extraction, and mapping these to an existing semantic structure such as a wordnet, can help us construct a valuable domain-specific semantic resource for any language and with minimum manual effort. However, in order to map the extracted terms in the target language onto a wordnet, we need a bilingual resource, preferably a domain-specific one, to provide the links between the source structure (in our case PWN) and the target structure (sloWNet). For our target domain of information science we have compiled a small parallel corpus of scientific abstracts and combined it with a bilingual online dictionary of computer science. Since both of these bilingual resources are used primarily to translate the hypernyms of the extracted terms, the parallel corpus does not need to fulfill all the requirements of a representative corpus.




Figure 1: Resources for harvesting MWEs

### 3.3 Automatic Term Extraction

The domain-specific Ikorpus is composed of texts from five journals dealing with computer science, information and communication technology, and it also con-

tains five consecutive volumes of proceedings of the largest informatics conference in Slovenia (DSI). Its size is approximately 15 million running words, which makes it an excellent and fairly representative source of terminology.

The task of automatically identifying domain-specific terms in texts has been addressed by numefrom PWN 2.0 into Slovenian authors and has been tackled to the extent that there now exist several commercial tools with term extraction functionality. The main approaches described in literature range from statistical ones, where terms are viewed as kinds of collocations, and the challenge lies in identifying the optimal word dependency measure (Dunning 1993; Daille 1995), to more linguistically informed and hybrid approaches, where part-of-speech, morphology and syntax are exploited as indicators of termhood (Heid 1999; Dias et al. 2000). More recent approaches introduce semantics and utilize context features to detect terminologically relevant phrases in running text (Maynard & Ananiadou 1999; Gillam, Tariq & Khurshid 2007), as well as propose methods for the identification of term variants (Jacquemin 2001). An overview of the trends is given in Kageura et al. (2004).

In our experiment, automatic term extraction is performed using a hybrid approach based on morphosyntactic patterns for Slovene and statistical ranking of candidates (Vintar 2004; 2009). The patterns, such as Adjective+Noun or Noun+Noun[Gen], yield numerous potential MWEs. After candidate phrases are extracted from the corpora, a term weighting measure is used to assign a “termhood” value to each phrase. The termhood value  $W$  of a candidate term  $a$  consisting of  $n$  words is computed as

$$W(a) = \frac{f_a^2}{n} * \sum_1^n \left( \log \frac{f_{n,D}}{N_D} - \log \frac{f_{n,R}}{N_R} \right)$$

where  $f_a$  is the absolute frequency of the candidate term in the domain-specific corpus,  $f_{n,D}$  and  $f_{n,R}$  are the frequencies of each constituent word in the domain-specific and the general language reference corpus respectively and  $N_D$  and  $N_R$  are the sizes of these two corpora in tokens.

The rationale of the termhood measure is that terms are composed of terminologically relevant words, and the measure of terminological relevance is the comparison of a word’s frequency between a domain-specific corpus and a general language corpus. This intuitive notion was first exploited by Ahmad et al. (1992) and implemented in other term extraction tasks (Scott 1998; Heid et al. 2001), however mostly for single-word terms. We use a modified version of this idea by adjusting it to multi-word expressions and including the frequency of the entire expression to override non-terminological phrases occurring only once.



### 3 Enriching Slovene wordnet with domain-specific terms

For example, if we compare two phrases a and b, both occurring in the corpus of computer science texts, *spletni brskalnik* ‘web browser’, (517) and *kakovost izdelka* ‘product quality’ (74), using the 619-million-words FidaPlus corpus as the source of comparative frequencies we get the following result which indicates that the first phrase is terminologically more relevant than the second:

$$(1) \quad W(a) = 517^2/2 * (9.18 + 2.93) = 1618434.89$$

$$(2) \quad W(b) = 74^2/2 * (1.31 + 1.20) = 6872.38$$

The term extraction procedure performed on the 15-million-token Slovene corpus of computer science yielded over 70,000 term candidates of length up to 4 words (Table 3). Given this bulk we can safely assume that not all of them are really terms we would like to include in sloWNet. As it turns out, the candidates list contains a large number of named entities, such as names of software and hardware products, vendors and manufacturers. Since few of these names might be terminologically relevant, we excluded them from further processing. We also employed a frequency threshold and discarded all term candidates which occurred less than 5 times.

The extractor uses morphosyntactic patterns, therefore each multi-word term candidate, e.g. *domenski strežnik* ‘domain name server’, can be automatically assigned a headword (*strežnik* ‘server’) and we assume this to be the hypernym of the term candidate.

Table 3: Term candidates and their length in words

MWE size	Number of candidates
2 words (Adj+N, N+N, ...)	54,844
3 words (Adj + Adj + N, N + Prep + N, ...)	16,861
4 words (Adj + Adj + N + N, ...)	2,605
<b>Total</b>	<b>74,310</b>

Clearly, the domain-specific terms constitute a valuable lexical resource, but not until we can introduce some semantic structure. The next step therefore is to integrate at least some of these terms into sloWNet.

## 4 Mapping terms to sloWNet

At this point we have a large number of Slovene multi-word terms without any semantic information other than the headword of each unit. Thus, for a term such as *prosto programje* ‘free software’, since it has been extracted through the syntactic pattern Adjective + Noun, we know that *programje* ‘software’ is the headword and *prosto* the modifier. We may also assume that *programje* is the hypernym of *prosto programje*, and hence we could add *prosto programje* into sloWNet as the hyponym of *programje*, but only if sloWNet already contains the required headword *programje*.

For many multi-word terms this turns out not to be the case, which is why we wish to add both the hypernym and its extracted hyponyms to sloWNet in order to fill as many lexical gaps as possible. We use the Princeton WordNet (PWN) as the source of semantic structure, and to be able to link headwords to this structure we use bilingual lexicon extraction.

### 4.1 Bilingual lexicon extraction

Bilingual lexicon extraction, also known as word alignment, is a statistical procedure where for each source word  $a$  the algorithm computes the probabilities of all of its potential translation equivalents  $t_1, t_2, t_n$  in the target language (Och & Ney 2003). The translation equivalents with the highest probability scores are then proposed as entries in the bilingual lexicon. Bilingual lexicon extraction can only be performed on parallel corpora or bitexts.

A small English-Slovene parallel corpus of 300,000 tokens is fed to the Uplug word aligner (Tiedemann 2003), which produces suggested translations for each word found in the corpus. To improve accuracy, we use only alignments of words that occur more than once and alignment scores over 0.05. This yields a bilingual single-word lexicon of 1326 words, mostly nouns (Table 4).

Table 4: Sample entries in the bilingual lexicon

Freq	Score	English	POS	Slovene	POS
4	0.058264988	adaptability	n	prilagodljivost	n
8	0.100445189	additional	a	dodaten	a
5	0.138443460	agent	n	agent	n

In order to improve coverage and accuracy, the automatically extracted bilingual lexicon is further enlarged with entries from the English-Slovene online dictionary of computer science Islovar. The dictionary provides approximately 5,000 bilingual entries and is consulted also in certain cases of ambiguous headword, as described below.

## 4.2 Adding Terms to sloWNet

For each Slovene multi-word term candidate we first identify its headword and assume that the headword is its hypernym. Using our bilingual lexicon we translate the headword into English and retrieve its synset IDs from PWN. If the headword turns out to be monosemous, the entire group of multi-word terms with the same hypernym can be added to sloWNet under the unique synset ID (Table 5).

Table 5: Monosemous headword

Term candidates	Hypernym, English translation and possible synset IDs	Selected synset ID
prosto programje ('free software')	programje = software	ENG20-06162514-n
priloženo programje ('attached software')		
ustrezno programje ('appropriate software')	ENG20-06162514-n	
novejše programje ('updated software')	computer_science	
dotatno programje ('additional software')		
vohunsko programje ('spyware')		

If the headword could be assigned several possible senses, we exploit the domain label in the wordnet, such as *factotum*, *biology* etc. If one of the senses of the polysemous headword belongs to the domain Computer Science, then this sense is chosen (Table 6).

If the headword is already part of sloWNet, no disambiguation is needed and the terms can be simply added as hyponyms to the existing Slovene hypernym. Also, in some cases one of the extracted multi-word terms was already in the Islovar dictionary. We can then use the English translation of the term to look up the correct hypernym and synset ID in PWN.

Nevertheless there remain many cases where the polysemous headword does not belong to the CompSci domain in the wordnet and it is neither included in sloWNet or Islovar. In such cases the correct sense must be picked manually (Table 7).

Table 6: Polysemous headword with CompSci domain

Term candidates	Hypernym, English translation and possible synset IDs	Selected synset ID
vgrajena tipkovnica ('built-in keyboard')	tipkovnica = keyboard	ENG20-03480332-n
brežžična tipkovnica ('wireless keyboard')		
zaslonska tipkovnica ('monitor keyboard')	ENG20-03480198-n	
tipkovnica qwerty ('QWERTY keyboard')	computer_science	
navidezna tipkovnica ('virtual keyboard')		
miniaturna tipkovnica ('miniature keyboard')	ENG20-03480332-n	
zunanja tipkovnica ('external keyboard')	factotum	
zložljiva tipkovnica ('folding keyboard')		
ergonomska tipkovnica ('ergonomic keyboard')		
programska tipkovnica ('program keyboard')		
slovenska tipkovnica ('Slovene keyboard')		
modularna tipkovnica ('modular keyboard')		
alfanumerična tipkovnica ('alphanumeric keyboard')		

Table 7: Polysemous headword, ID to be selected manually

Term candidates	Hypernym, English translation and possible synset IDs	Selected synset ID
nalaganje gonilnikov ('loading drivers')	nalaganje = loading	to be selected manually
nalaganje podatkov ('loading data')		
nalaganje programov ('software download')	ENG20-00671518-n	
nalaganje strani ('loading page')	factotum	
	ENG20-13044298-n	
	transport	

## 5 Discussion

Extracting terms from a large domain-specific Slovene corpus yields the bulk of 74,310 term candidates. We keep only those that occur more than five times and where the headword and its English translation can be identified with reasonable accuracy, and we disregard all names and terms that include names. Some of the remaining terms were already either in the Islovar dictionary or in sloWNet, however the large majority were new. Table 8 shows the number of terms successfully added to sloWNet.

The assumption that the headword of the multi-word expression is at the same time the hypernym of the term may seem daring, however we encountered very

### 3 Enriching Slovene wordnet with domain-specific terms

Table 8: Total term candidates added to sloWNet

Category	Number of terms
Already in sloWNet	29
Already in PWN	23
Already in Islovar	198
New	5150
<b>Total</b>	<b>5400</b>

few examples where this is not the case. Within a random sample of 200 multi-word terms we found 5 terms where the headword could not be considered an appropriate hypernym of the term, for example a *spletni portal* ‘web portal’ is not a kind of *portal* ‘portal’; *portal* being an architectural term, and *prostor na disku* ‘disk space’ is not a kind of *prostor* ‘space’; although both of these headwords could be used elliptically in a computer science context to mean *web portal* or *disk space* respectively.

As has been described in the previous section, the difficult part is determining the correct sense of the potentially polysemous headword. This ambiguity can of course affect a large number of terms, since – as can be seen in Table 6 – several dozens of multi-word terms share the same headword. While we use all the semantic information we can infer either from the domain label or the online dictionary, nearly half of all the headwords need to be disambiguated manually (Table 9).

Table 9: Categories of headwords

Category	Number of headwords
Monosemous	84
Headwords with CompSci domain	35
Headwords already in sloWNet	11
Headwords derived from MWE PWN	6
To be picked manually	136
<b>Total</b>	<b>272</b>

In this respect our methodology could benefit significantly from additional context-based disambiguation procedures. A possible approach would be to use

the contexts of the polysemous headwords and compute the semantic similarity between the relevant context words and each sense of the headword. The sense with the greatest semantic similarity to the context features is selected as the correct one. This is essentially a word disambiguation task and various authors have proposed similarity measures based on the graph representation of wordnets (e.g. Leacock & Chodorow 1998; Wu & Palmer 1994; Agirre et al. 2009). In future experiments we plan to implement such methods for the selection of the correct sense.

Finally it should be noted that the domain labels in Princeton WordNet are sometimes illogical, too specific or not specific enough. If we for example explore the financial domain, there is no label [Finance], but we find three different domains for a related set of concepts: *money* [Money], *coin* [Money], *bank* [Banking], *account* [Banking], *pay* [Economy]. This is clearly a problem for automatic text processing, because we cannot rely on the fact that semantically related lexical items share the same domain label in WordNet. On the other hand there exists a hierarchical structure of WordNet domains which was not taken into account in our experiments. It may be the case that some ambiguity issues could be better resolved using this hierarchy.

## 6 Conclusions

We described an approach to improve the domain coverage of a wordnet by enriching it with semi-automatically extracted multi-word terms. Our method utilizes a combination of mono- and bilingual resources. A large monolingual domain-specific corpus is used as the source of terminology, and a smaller parallel corpus combined with a domain-specific dictionary is used to provide translation equivalents of headwords. These are required in order to map the semantic structure of Princeton WordNet onto the Slovene term candidates and thus integrate them into sloWNet.

Although the approach works well and yields many items of specialised vocabulary, the most difficult part is the selection of the correct sense with polysemous headwords. In some cases the correct sense can be inferred from the domain label or from the dictionary, but in many cases this step still has to be performed manually. In the future we plan to implement a sense disambiguation procedure based on semantic similarity.

It should be noted that an evaluation of monolingual term extraction lies beyond the scope of this paper and is not addressed, although the quality of the term candidates clearly influences the results of the experiment described. Term ex-

traction evaluation depends heavily on the target application, which means that the same system may perform very well in an information retrieval task and poorly in a dictionary-making task. Since the measure of terminological relevance relies on the comparison of relative frequencies between a domain-specific and a reference corpus, the term extraction system performs better for highly specialised domains or, in other words, for terms that do not occur frequently in general language. Information science is in this respect not the ideal domain because IT-related topics are regularly discussed in general language media.

The proposed methodology can be extended to other domains, or indeed other languages. While we employ a specialised monolingual corpus, a bilingual corpus and a specialised bilingual dictionary, the cross-language part of the algorithm is essentially suited to parallel corpora. Especially in domains – or language pairs – for which bilingual dictionaries are scarce it is often more viable to construct a small parallel corpus and use the word-aligned bilingual lexicon to translate headwords. While in other domains we could again exploit the domain labels in WordNet to disambiguate the headword, our methodology is less suitable for general language where polysemy is common and disambiguation can only be performed with context-based methods.

An evaluation of the domain coverage of sloWNet will be performed within a Machine Translation application. In the future we also plan to extend this approach to the extraction of definitions from domain-specific corpora using Machine Learning to distinguish between well-formed and not-well-formed definitions (Fišer, Pollak & Vintar 2010).

## References

- Agirre, Eneko, Enrique Alfonseca, Keith Hall, AJana Kravalova, Marius Pasca & Aitor Soroa. 2009. A study on similarity and relatedness using distributional and WordNet-based approaches. In *Proceedings of NAACL-HLT 09*. Boulder, USA.
- Ahmad, Khurshid, Andrea Davies, Heather Fulford & Margaret Rogers. 1992. What is a term? The semi-automatic extraction of terms from text. In Mary Snell-Hornby, Franz Pöchhacker & Klaus Kaindl (eds.), *Translation studies – an interdisciplinary*, 267–278. Amsterdam & Philadelphia: John Benjamins.
- Bourigault, Didier. 1993. Analyse syntaxique locale pour le repérage de termes complexes dans un texte. *Traitement Automatique des Langues* 34(2). 105–117.
- Daille, Béatrice. 1995. Combined approach for terminology extraction: Lexical statistics and linguistic filtering. In *COLING 94*, 515–521.

- Diab, Mona & Philip Resnik. 2002. An unsupervised method for word sense tagging using parallel corpora. In *40th Anniversary Meeting of the Association for Computational Linguistics (ACL-02)*, 255–262. Philadelphia: Association for Computational Linguistics.
- Dias, Gaël & Sérgio Nunes. 2004. Evaluation of different similarity measures for the extraction of multiword units in a reinforcement learning environment. In *Proceedings of the 4th International Conference on Languages Resources and Evaluation*, 1717–1721. Lisbon, Portugal.
- Dias, Gaël, Sylvie Guilloré, Jean Claude Bassano & José Gabriel Pereira Lopes. 2000. Combining linguistics with statistics for Multiword Term Extraction: A fruitful association? In *Proceedings of 6ème Conférence sur la recherche d'informations assistée par ordinateur (RIAO 2000)*, 1–20. Paris, France.
- Dunning, Ted. 1993. Accurate methods for the statistics of surprise and coincidence. *Computational linguistics* 19. 61–74.
- Dyvik, Helge. 1998. Translations as semantic mirrors. In *Proceedings of Workshop W13: Multilinguality in the lexicon ii. The 13th biennial European Conference on Artificial Intelligence ECAI 98*, 24–44. Brighton, UK.
- Erjavec, Tomaž & Darja Fišer. 2006. Building slovene wordnet. In *Proceedings of the 5th International Conference on Language Resources and evaluation LREC'06*. Genoa, Italy. <http://www.lrec-conf.org/lrec2006/nl.ijs.si/et/Bib/LREC06/lrec06-slownet-150.pdf>.
- Erjavec, Tomaž & Nancy Ide. 1998. The MULTTEXT-East corpus. In *Proceedings of the First International Conference on Language Resources and Evaluation, LREC'98*. Granada, Spain.
- Fellbaum, Christiane. 1998. *WordNet: An electronic lexical database*. Cambridge: MIT Press.
- Fišer, Darja. 2007. Leveraging parallel corpora and existing wordnets for automatic construction of the Slovene wordnet. In *Proceedings of the 3rd Language and Technology Conference L& TC'07*. Poznan, Poland.
- Fišer, Darja, Senja Pollak & Špela Vintar. 2010. Learning to mine definitions from Slovene structured and unstructured knowledge-rich resources. In *Proceedings of the 7th Language Resources and Evaluation Conference*. Malta.
- Fišer, Darja & Benoît Sagot. 2008. Combining multiple resources to build reliable wordnets. In *Text, Speech and Dialogue Conference (LNCS 2546)*, 61–68. Berlin & Heidelberg: Springer.
- Gillam, Lee, Mariam Tariq & Ahmad Khurshid. 2007. Terminology and the construction of ontology. In Fidelia Ibekwe-SanJuan, Anne Condamines & M.



### 3 Enriching Slovene wordnet with domain-specific terms

- Teresa Cabré Castellví (eds.), *Application-driven terminology engineering*, 49–74. Amsterdam & Philadelphia: John Benjamins.
- Heid, Ulrich. 1999. Extracting terminologically relevant collocations from German technical texts. In Peter Sandrini (ed.), *Terminology and knowledge engineering (TKE99)*, 241–255. Vienna: TermNet.
- Heid, Ulrich, Stefan Evert, Arne Fitschen, Marion Freese & Andreas Vögele. 2001. *Term candidate extraction in DOT: Dot final report, Part II*. Stuttgart: IMS, University of Stuttgart.
- Horák, Aleš, Karel Pala, Adam Rambousek & Martin Povolný. 2005. DEBVisDic – first version of new client-server wordnet browsing and editing tool. In *Proceedings of the Global Wordnet Conference GWA'05*, 325–328. Brno.
- Ide, Nancy, Tomaž Erjavec & Dan Tufis. 2002. Sense discrimination with parallel corpora. In *Proceedings of ACL'02 Workshop on Word Sense Disambiguation: Recent successes and future directions*, 54–60. Philadelphia.
- Jacquemin, Christian. 2001. *Spotting and discovering terms through Natural Language Processing*. Cambridge, Massachusetts: MIT Press.
- Kageura, Kyo, Béatrice Daille, Hiroshi Nakagawa & Lee-Feng Chien. 2004. Introduction: Recent trends in computational terminology. *Terminology* 10(1). 1–21.
- Krstev, Cvetana, Gordana Pavlović-Lažetić, Duško Vitas & Ivan Obradović. 2004. Using textual resources in developing Serbian wordnet. *Romanian Journal of Information Science and Technology* 7(1–2). 147–161.
- Leacock, Claudia & Martin Chodorow. 1998. Combining local context and WordNet sense similarity for word sense identification. In Christian Fellbaum (ed.), *WordNet: An electronic lexical database*, 265–283. Cambridge: MIT Press.
- Maynard, Diana & Sophia Ananiadou. 1999. Term extraction using a similarity-based approach. In Didier Bourigault, Christian Jacquemin & Marie-Claude L'Homme (eds.), *Recent advances in computational terminology*, 261–278. Amsterdam: John Benjamins.
- Och, Franz-Josef & Hermann Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics* 29(1). 19–51.
- Piao, Scott, Paul Rayson, Dawn Archer, Wilson Andrew & Tony McEnery. 2003. Extracting multiword expressions with a semantic tagger. In *Workshop on multiword expressions of the 41st ACL meeting*, 49–57. Sapporo, Japan.
- Scott, Mike. 1998. Focusing on the text and its key words. In *TALC 98 proceedings*, 152–164. Oxford.
- Steinberger, Ralf, Bruno Pouliquen, Anna Widiger, Camelia Ignat, Tomaž Erjavec, Dan Tufiş & Dániel Varga. 2006. The JRC-Acquis: A multilingual aligned paral-

- lel corpus with 20+ languages. In *Proceedings of the 5th International Conference on Language Resources and Evaluation*. Genoa, Italy.
- Tiedemann, Jörg. 2003. *Recycling translations: Extraction of lexical data from parallel corpora and their application in natural language processing*. Doctoral Thesis. *Studia Linguistica Upsaliensia* 1 PhD thesis.
- Tomokiyo, Takashi & Matthew Hurst. 2003. A language model approach to keyphrase extraction. In *Workshop on Multiword Expressions of the 41st ACL meeting*, 33–41. Sapporo, Japan.
- Tufis, Dan. 2000. BalkaNet – design and development of a multilingual Balkan WordNet. *Romanian Journal of Information Science and Technology Special Issue* 7(1–2).
- Vintar, Špela. 2004. Comparative evaluation of C-value in the treatment of nested terms. In *Memura 2004 – Methodologies and evaluation of Multiword Units in Real-World Applications (LREC 2004)*, 54–57.
- Vintar, Špela. 2009. Samodejno luščenje terminologije – izkušnje in perspective. Automatic term recognition – experience and perspectives. In N. Ledinek, M. Žagar Karer & M. Humar (eds.), *Terminologija in sodobna terminografija*, 345–356. Ljubljana: Založba ZRC.
- Vintar, Špela & Darja Fišer. 2008. Harvesting multi-word expressions from parallel corpora. In *Proceedings of the 6th International Conference on Language Resources and Evaluation LREC'08*. Marrakech, Morocco.
- Wu, Zhibiao & Martha Palmer. 1994. Verb semantics and lexical selection. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*. Las Cruces, New Mexico.