ISSN 2007-9737

Discovering Continuous Multi-word Expressions in Czech

Zuzana Nevěřilová

Masaryk University, Faculty of Informatics, Brno,
Czech Republic

xpopelk@fi.muni.cz

Abstract. Multi-word expressions frequently cause incorrect annotations in corpora, since they often contain foreign words or syntactic anomalies. In case of foreign material, the annotation quality depends on whether the correct language of the sequence is detected. In case of inter-lingual homographs, this problem becomes difficult. In the previous work, we created a dataset of Czech continuous multi-word expressions (MWEs). The candidates were discovered automatically from Czech web corpus considering their orthographic variability. The candidates were classified and annotated manually. Afterwards, the dataset was extended automatically by generating all word forms of those MWEs that were annotated as nouns. In this work, we used the dataset as positive examples, we filtered out negative examples from the MWE candidates. We trained a classifier with mean accuracy 92.7%. We have shown that the combined approach slightly outperforms approaches concerning only association measures mainly on MWEs containing inter-lingual homographs and out-of-vocabulary words. The discovery methods can be applied to other languages which encounter orthographic variability in web corpora.

Keywords. Multiword expression, multi-word expression, MWE, MWE discovery, inter-lingual homographs

1 Introduction

Multi-word expressions (MWEs), consist of several words but behave as a single word to some extent [5]. Their idiosyncracy causes (among others) problems in corpus annotation which is conventionally token-based.

MWEs do not form a homogeneous group. [5] point out four main characteristics of MWEs: syntactic anomaly, non-compositionality, non-substitutability, and ambiguity between MWE and non-MWE readings. It is nevertheless necessary to mention that not all MWEs have all four characteristics.

In [15], a taxonomy of MWEs consists of two basic groups: lexicalized phrases and institutionalized phrases. The former group is broken down into three subgroups: fixed expressions, semi-fixed expressions, and syntactically-flexible expressions. Fixed MWEs never change and are sometimes considered as a word with spaces, e.g. *ad hoc*. Semi-fixed expressions "undergo some degree of lexical variation, e.g. in the form of inflection, variation in reflexive form, and determiner selection". [15]. Syntactically-flexible expressions have a larger degree of syntactic variability including word order and possible gaps.

In this work, we focus on fixed and semi-fixed expressions and their discovery. The aim was to create an extensive list of MWEs that will be used for Czech corpora annotation. We show that traditional methods based on association measures work well for a considerable number of MWEs, however a large subclass of MWEs stay undiscovered. We therefore proposed a method based on discovery of orthographic variability. Using this method, we extracted 26,704 MWE candidates that were annotated manually by four annotators who marked about 5,800 of the candidates as MWEs.

We observed other features of the annotated data and we built a classifier that can be used for discovery of further MWEs.The paper is organized as follows: Section 2 introduces MWE annotation, Section 3 references MWE discovery methods and MWE-annotated corpora. In Section 4, we describe in detail the construction of the MWE

dataset. Section 5 summarizes results of this work. A concluding discussion is found in Section 6.

2 MWE Annotation

Annotation pipelines mostly consist of tokenization, morphological analysis, and tagging. A naive approach would be to create a large list and to treat MWEs as words with spaces, i.e. to tokenize sentences like *It is a priori impossible* to (*It*, *is*, *a priori*, *impossible*).

This approach is suitable for fixed, non-decomposable expressions, less suitable for semi-fixed expressions, and unsuitable for syntactically-flexible expressions.

For example, [15] shows that the verb-particle construction *look up* can have two meanings (to look upwards and to search) in some contexts and only one in others. Lists are not sufficiently *flexible* to cope with this ambiguity.

Moreover, MWE lists should be extensive yet not complete. This lack of generality is called *lexical proliferation problem* in [15].

The approaches such as [9, 21] annotate MWEs both as single tokens and as MWEs. For example, in Wiki50, MWEs are split into tokens and annotated according to the Inside Outside Beginning (IOB) standard [1]. The corpus LexSem [17] distinguishes strong and weak multi-word groupings.

2.1 MWE Annotation in Czech Corpora

The pipelines for Czech corpus annotation do not consider MWEs at all, at least in case of the web corpus `czTenTen` [18] and SYNv6 corpus provided by the Czech National Corpus [7]. Instead, parts of the most frequent MWEs were included into the dictionaries used by morphological analyzers.

For example, the word *priori* (being part of the MWE *a priori*) is in morphological dictionaries annotated as an adverb. As a result, *a priori* is annotated as two tokens, one being (incorrectly) a conjunction, another being an adverb. Apart from the Latin preposition, *a* is also a Czech conjunction (meaning *and*).

Similarly, the fixed MWE *hot dog* is annotated as an unknown token *hot* and a noun *doga* (a dog breed *Great Dane*. In `czTenTen`, *hot* is annotated as an interjection. These two examples illustrate what problems in MWE annotation the out-of-vocabulary words and inter-lingual homographs cause.

3 Related Work

Multi-word expressions appear in a wide range of NLP tasks such as machine translation (e.g. [16]), parsing [6], or lexicography (e.g. [10]). Therefore the number of works concerning MWEs is high. In the following text, we focus on MWE discovery as well as on works concerning Czech.

3.1 MWE Discovery

Early discovery approaches worked with collocation measures. Association measures, namely point-wise mutual information (first proposed by [4]) are among the most popular discovery methods. Later, works as [13] show that adding linguistic information improves the results.

Another group of approaches is based on lexical fixedness. For example, [20] use K-means clustering algorithm with cosine similarity to build an unsupervised approach to MWEs discovery.

MWEs are also discovered employing semantic properties: e.g. [8] use latent semantic analysis to identify non-compositional MWEs. An example approach combining several information sources is [19]. In this work, authors use among others orthographic variations with hyphens.

3.2 MWE Annotated Corpora

Although there are shared tasks at SIGLEX and working groups in PARSEME, not many MWE annotated corpora exist. Examples of such corpora are the social web corpus [17], French corpus with annotated multiword nouns [9], or the corpus of 50 Wikipedia articles with annotated MWEs – Wiki50 [21], or parallel corpus of TED talks [11].

3.3 MWEs in Czech

Czech MWEs are studied mainly with respect to their syntactic structure. SemLex, the lexicon of Czech MWEs, is used for syntactic identification of MWE occurrences in text. This process is described in detail in [3]. SemLex was built by identifying MWEs in the Prague Dependency Treebank [2].

4 Building the Dataset for Czech MWE Annotation

In this Section, we describe the process of automatic discovery of MWE candidates, manual classification, comparison with association measures, and training. We explain different aspects of the manual annotation that influence selection of the training examples.

4.1 MWE Discovery Based on Orthographic Variability

We made several observations on Czech MWEs and we found fluctuating orthography of frozen MWEs in the Czech web corpus `czTenTen` [18]. In other words, people are sometimes unsure whether the correct Czech orthography for expressions such as *a priori* is *apriori* or even *a-priori* (the correct variants are *a priori* and *apriori*). Similar experience is mentioned in [9] and [19] for different languages.

Using a web corpus was a key decision: first, `czTenTen` is so far one of the largest Czech corpora, second, web corpus contains many kinds of (mostly unedited) texts. Our observations indicated that e.g. in discussion groups, orthographic variants appear frequently since people use a language between correct written Czech and spoken Czech and do not care much about the correctness.

The MWE discovery is described more in detail in [12]. The approach is straightforward: if a chunk exists in corpus in all three forms (several words, one word, several words with dashes) above a minimum frequency threshold, we consider it a MWE candidate. This method discovered 26,704 MWE candidates.

4.2 Classification of MWE Candidates and Observation

The classification of MWEs was manual: four annotators had to decide whether a sequence of words is random (non-MWE), MWE with function of a noun, MWE with function of an adverb, unspecified English loanword or other foreign. The resulting collection contained 3,219 MWEs with function of a noun, 80 MWEs with function of an adverb, 2,325 English and 140 non-English foreign MWEs.

There were only 33 candidates that were annotated differently by the four annotators. We included all MWE candidates with a majority agreement in the dataset.

A more detailed observation of the classified data sample has shown that the positive entries (MWEs of any kind) do not contain evident non-MWEs. The high inter-annotator agreement was also caused by the similarity of the annotators: they shared the same field of study, interests, and age. On the other hand, observation of negative entries sample indicated that the dataset coverage is limited. We found two main reasons why the negative entries were rather noisy: First, many actual MWEs were annotated as non-MWE since the annotators did not understand the expression. Second, annotators sometimes did not annotate MWEs with non-standard spelling. Some MWEs are frequent with incorrect spelling, for example *á propos* has 876 occurrences in `czTenTen` while the correct spelling *à propos* has only 147 occurrences[1].

After observing the positive entries we could distinguish several types of Czech MWEs:

— **Czech fixed phrases** with syntactic anomalies (e.g. *stůj co stůj*, lit. imperative *go what go* meaning *by hook or crook*.

— **Non-English borrowings** which are not analyzable by most users of the Czech language (e.g. *faux pas*, *a priori*).

[1] The main reason in this case could be that the character *à* is not in the Czech keyboard layout.

— **English calques** (loan translations) that are syntactically anomalous in Czech (e.g. *risk management* contains a noun modifier which precedes the head noun – such construction does not originally exist in Czech)

— **Proper names** (e.g. San Francisco, Air France). Although [15] show that location names in sport club names are ellidable (e.g. *the (San Francisco) 49ers*), proper names of sport clubs, companies, locations are still highly idiosyncratic. Personal names do not share this fixedness level, mainly for two reasons: first, they are not strictly continuous, since we can find constructions such as *Barack Obambi Obama*, second, most personal names (e.g. *John Smith*) are tuples formed from lists of first names (e.g. *John*, *Jane*) and last names (e.g. *Smith*). They are therefore decomposable and their components can be combined (e.g. *Jane Smith*).

The first two categories contain fixed MWEs whereas the third and fourth categories cover words that are often subject to inflection (semi-fixed MWEs). In some cases, the inflectional pattern is not clear to language users so they avoid inflection. For example, users prefer expressions such as *jít do obchodu Marks & Spencer* (*go to the store Marks & Spencer*) over (shorter) *jít do Markse & Spencera* (*go to Marks & Spencer's*).

The correct inflection requires gender assignment which is based mostly on the word ending. Roughly said, language users assign masculine inanimate gender to words ending with a consonant, feminine gender to words ending with *-a* or *-e* and neuter gender to words ending with *-o*. This rule is sometimes influenced by similar words that were adopted previously. For example, *after party* has the same gender as *party* (feminine) since *party* exists in Czech for decades. On the other hand, the gender of *Air France* is unclear and probably it is not assigned at all.

4.3 Automatic Extension of the Dataset

As mentioned above, semi-fixed MWEs are subject of inflection which is quite regular in Czech.

We decided to extend the dataset with automatically generated word forms.

Within the manual annotations, we observed that annotators were often unsure whether English calques such as *news server* have to be annotated as nouns or unspecified English phrases. This indistinctness disappeared in case of inflected MWEs: e.g. if the annotator had to decide whether *news serveru* (genitive) is a noun or an unspecified English phrase, she was sure that it is a (Czech) noun. Therefore, we decided to annotate all English calques such as *news server* as nouns if we found corpus evidence for their inflection, i.e. at least one inflected form.

Eventually, we modified the annotation for MWEs with the same structure containing the same head noun. For example, we observed that *news server* is subject of inflection therefore we also inflected *mail server*, *DHCP server* etc. This modification was controlled manually since the rule is not generally applicable. For example, in case of *client server* the inflection makes no sense since the similarity in structure is only shallow (i.e. *client* is not modifying *server*).

From nominative singular forms, we generated genitive, dative, accusative, locative, and instrumental. We also generated plural forms for those MWEs that are not named entities. We did not generate vocative since it is used only for animate nouns. Also, the plural forms of named entities are rather rare, so we did not generate them. As a result, we obtained a dataset of 24,807 word forms.

4.4 Preparing Training Data

Next step in the task was to examine the relationship between traditional association measures and the annotated data. According to [5], association measures are straightforward for two-word expressions but rather complicated for longer ones. In the dataset, 24,360 entries are two-word expressions, 423 entries are three-word expressions, and 24 are longer. We decided not to take longer entries into account since they make only 1,8% of the data.

In Section 4.2, we described the nature of positive (MWEs) and negative (non-MWEs) annotations.

Although the majority of the candidates mentioned in Section 4.1 were annotated as non-MWEs, we could not simply use them as negative examples. We filtered out named entities automatically and made further manual cleaning.

T-score:

$$\frac{f_{xy} - \frac{f_x f_y}{N}}{\sqrt{f_{xy}}}, \tag{1}$$

MI-score:

$$\log_2 \frac{f_{xy} N}{f_x f_y}, \tag{2}$$

MI^3-score:

$$\log_2 \frac{f_{xy}^3 N}{f_x f_y}, \tag{3}$$

min. sensitivity:

$$\min(\frac{f_{xy}}{f_y}, \frac{f_{xy}}{f_x}), \tag{4}$$

logDice:

$$14 + \log_2 \frac{2 \cdot f_{xy}}{f_x + f_y}, \tag{5}$$

log-likelihood:

$$\begin{aligned} &2 \cdot (xlx(f_{xy}) + xlx(f_x - f_{xy}) + \\ &xlx(f_y - f_{xy}) + xlx(N) + xlx(N + f_{xy} - f_x - f_y) \\ &\quad - xlx(f_y) - xlx(N - f_x) - xlx(N - f_y)), \end{aligned} \tag{6}$$

where $xlx = f \ln(f)$.

In order to avoid skewed classes, we selected randomly the same number of positive and negative examples.

We computed association measures T-score, MI-score, MI^3-score, minimum sensitivity, logDice, and log-likelihood as described in [14]. Before applying the formulas above, the data were rescaled to $[-1, 1]$ using the function:

$$r = \frac{(2x - max - min)}{(max - min)}.$$

We are aware that the association measures are not independent features, therefore we employed the recursive feature elimination (RFE) in order to suppress less effective features. The best results (i.e. matching the largest number of positive examples and not matching the largest number of negative examples) were provided by MI-score, MI^3 and logDice. Roughly said, if the MI-score is a large number, the bigram is likely to be a MWE.

In Section 1, we mentioned that association measures do not work well for a large group of MWEs. For MWEs containing inter-lingual homographs, the MI-score is somewhat lower than for MWEs without inter-lingual homographs but still higher than for random bigrams. However, the minimum sensitivity is significantly lower for MWEs with homographs than for MWEs without homographs. This discrepancy which can be seen in Figure 1 causes problems in MWE discovery. MWEs containing inter-lingual homographs are often not discovered by means of association measures.

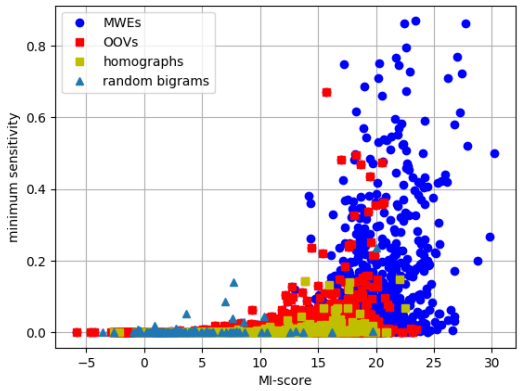

Fig. 1. MWEs have often high MI-score and higher minimum sensitivity than non-MWEs. However, MWEs containing inter-lingual homographs and OOVs have lower MI-scores and minimum sensitivity than other MWEs

We decided to include the information about inter-lingual homography as a feature using a list of 251 Czech-English homographs. Similarly, we added the information whether the bigram contains OOV words.

4.5 Classifier Training

We used logistic regression on 5,792 examples. Using 4-fold cross-validation, the classifier had 91.4% mean accuracy without using information about OOV and homography as a feature.

The information about homography increased the mean accuracy to 91.8%. Information about OOV increased the mean accuracy up to 92.7%. Eventually, information about OOV proved to be more useful.

5 Results

We obtained two results: a dataset of fixed and semi-fixed MWEs, and a classifier for discovering MWEs among bigrams.

Currently, the resource contains 4,731 MWE lemmata (24,807 word forms). The Table 1 shows different categories of the entries. Most of the MWEs have function of a noun, 634 are indeclinable.

We compared our dataset to SemLex [3], a manually classified list of 12,233 collocations found in Prague Dependency Treebank by three annotators. Here the 9,660 non-MWEs and 2,572 MWEs are classified as:

— non-collocations,

— stock phrases,

— proper names,

— support verb constructions,

— technical terms,

— idiomatic expressions.

The overlap in both resources is very small, only 40 lemmata are in common. The reason can be different view on what is a MWE and also different source data: Prague Dependency Treebank contains mostly correct Czech sentences, while web corpora often contain non-standard language and sequences with language mixing. In our resource, we focused on cases where incorrect annotation is more likely: non-standard, syntactically anomalous, containing OOVs or interlingual homographs.

The classifier was trained on 75% of the example data and tested on 25%. In the 4-fold evaluation, the resulting mean accuracy is 0.93, precision 0.94, recall 0.96, and F1-score 0.95. We also measured the classifier on the SemLex data.

Table 1. Overview of Czech MWEs dataset

category	# of entries	# of lemmata
foreign	2,221	2,221
nouns	22,459	2,422
adjectives	42	3
adverbs	81	81
particles	4	4

In this case, the precision was 0.65, recall 0.66, and F1-score 0.65.

6 Conclusion and Future Work

Frozen continuous MWEs are in many cases incorrectly annotated in Czech corpora. It is caused by the idiosyncratic nature of such MWEs: they often contain rare of foreign words and they sometimes evince syntactic anomalies. The aim of this work is to discover MWEs.

The paper presents a new dataset of Czech fixed and semi-fixed MWEs. We described the acquisition of the data and the annotation process. The annotated data were automatically extended since many MWEs are subject to inflection. Finally, we used the data for classifier training.

To our knowledge the dataset is the largest list of fixed and semi-fixed MWEs. Another dataset for Czech, SemLex contains all types of MWEs including syntactically-flexible ones. The overlap with SemLex [3], is insignificant: only 40 MWE lemmata occur in both resources. The results are difficult to compare with other works: some deal with syntactically-flexible MWEs which are more difficult to discover, some work for languages with weak inflection.

We plan to use the dataset and the classifier for identifying MWEs in new version of the Czech web corpus `czTenTen`. This application can provide extrinsic evaluation if a measure of quality of corpus annotation will be defined.

Acknowledgements

This work has been partly supported by the Ministry of Education of CR within the OP VVV project CZ.02.1.01/0.0/0.0/16_013/0001781.

References

1. **Baldwin, B. (2009).** Coding Chunkers as Taggers: IO, BIO, BMEWO, and BMEWO+. [accessed 2017-09-28].
2. **Bejček, E., Hajičová, E., Hajič, J., Jínová, P., Kettnerová, V., Kolářová, V., Mikulová, M., Mírovský, J., Nedoluzhko, A., Panevová, J., Poláková, L., Ševčíková, M., Štěpánek, J., & Zikánová, Š. (2013).** Prague Dependency Treebank 3.0. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics, Charles University.
3. **Bejček, E., Straňák, P., & Pecina, P. (2013).** Syntactic Identification of Occurrences of Multiword Expressions in Text using a Lexicon with Dependency Structures. *Proceedings of the 9th Workshop on Multiword Expressions*, Association for Computational Linguistics, Atlanta, Georgia, USA, pp. 106–115.
4. **Church, K. W. & Hanks, P. (1990).** Word Association Norms, Mutual Information, and Lexicography. *Comput. Linguist.*, Vol. 16, No. 1, pp. 22–29.
5. **Constant, M., Eryiğit, G., Monti, J., van der Plas, L., Ramisch, C., Rosner, M., & Todirascu, A. (2017).** Multiword Expression Processing: A Survey. *Computational Linguistics*, Vol. 0, No. ja, pp. 1–92.
6. **Eryiğit, G., İlbay, T., & Can, O. A. (2011).** Multiword Expressions in Statistical Dependency Parsing. *Proceedings of the Second Workshop on Statistical Parsing of Morphologically Rich Languages*, SPMRL '11, Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 45–55.
7. **Hnátková, M., Křen, M., Procházka, P., & Skoumalová, H. (2014).** The SYN-series Corpora of Written Czech. **Calzolari, N., Choukri, K., Declerck, T., Loftsson, H., Maegaard, B., Mariani, J., Moreno, A., Odijk, J., & Piperidis, S.**, editors, *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, European Language Resources Association (ELRA), Reykjavik, Iceland.
8. **Katz, G. & Giesbrecht, E. (2006).** Automatic Identification of Non-compositional Multi-word Expressions Using Latent Semantic Analysis. *Proceedings of the Workshop on Multiword Expressions: Identifying and Exploiting Underlying Properties*, MWE '06, Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 12–19.
9. **Laporte, E., Nakamura, T., & Voyatzi, S. (2008).** A French corpus annotated for multiword nouns. *Language Resources and Evaluation Conference. Workshop Towards a Shared Task on Multiword Expressions*, Marrakech, Morocco, pp. 27–30.
10. **Loukachevitch, N. & Lashevich, G. (2016).** Multiword expressions in Russian thesauri RuThes and RuWordnet. *2016 IEEE Artificial Intelligence and Natural Language Conference (AINL)*, pp. 1–6.
11. **Monti, J., Sangati, F., & Arcan, M.**, **.** TED-MWE: a bilingual parallel corpus with MWE annotation: Towards a methodology for annotating MWEs in parallel multilingual corpora. *Proceedings of the Second Italian Conference on Computational Linguistics CLiC-it 2015*, Accademia University Press, Torino.
12. **Nevěřilová, Z. (2015).** Annotation of Multi-Word Expressions in Czech Texts. **Horák, A., Rychlý, P., & Rambousek, A.**, editors, *Ninth Workshop on Recent Advances in Slavonic Natural Language Processing*, Tribun EU, Brno, pp. 103–112.
13. **Ramisch, C., Schreiner, P., Idiart, M., & Villavicencio, A. (2008).** An Evaluation of Methods for the Extraction of Multiword Expressions. *Proceedings of the LREC Workshop Towards a Shared Task for Multiword Expressions MWE 2008*, Marrakech, Morocco, pp. 50–53.
14. **Rychlý, P. (2008).** A Lexicographer-Friendly Association Score. **Horák, A. & Sojka, P.**, editors, *2th Workshop on Recent Advances in Slavonic Natural Language Processing*, Masaryk University, Brno, pp. 6–9.
15. **Sag, I. A., Baldwin, T., Bond, F., Copestake, A., & Flickinger, D. (2002).** Multiword Expressions: A Pain in the Neck for NLP. In **Gelbukh, A.**, editor, *Computational Linguistics and Intelligent Text Processing*, volume 2276 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, pp. 1–15.
16. **Sakamoto, S., Ogawa, Y., Nakamura, M., Ohno, T., & Toyama, K. (2017).** Utilization of Multi-word Expressions to Improve Statistical Machine Translation of Statutory Sentences. **Otake,**

M., Kurahashi, S., Ota, Y., Satoh, K., & Bekki, D., editors, *New Frontiers in Artificial Intelligence: JSAI-isAI 2015 Workshops, LENLS, JURISIN, AAA, HAT-MASH, TSDAA, ASD-HR, and SKL, Kanagawa, Japan, November 16-18, 2015, Revised Selected Papers*, Springer International Publishing, Cham, pp. 249–264.

17. **Schneider, N., Onuffer, S., Kazour, N., Danchik, E., Mordowanec, M. T., Conrad, H., & Smith, N. A. (2014).** Comprehensive Annotation of Multiword Expressions in a Social Web Corpus. **Calzolari, N., Choukri, K., Declerck, T., Loftsson, H., Maegaard, B., Mariani, J., Moreno, A., Odijk, J., & Piperidis, S.**, editors, *Proceedings of the Ninth International Conference on Language Resources and Evaluation*, ELRA, Reykjavík, Iceland, pp. 455–461.
18. **Suchomel, V. (2012).** Recent Czech Web Corpora. **Horák, A. & Rychlý, P.**, editors, *6th Workshop on Recent Advances in Slavonic Natural Language Processing*, Tribun EU, Brno, pp. 77–83.
19. **Tsvetkov, Y. & Wintner, S. (2011).** Identification of Multi-word Expressions by Combining Multiple Linguistic Information Sources. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP '11, Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 836–845.
20. **Van de Cruys, T. & Moirón, B. V. (2007).** Semantics-based Multiword Expression Extraction. *Proceedings of the Workshop on a Broader Perspective on Multiword Expressions*, MWE '07, Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 25–32.
21. **Vincze, V., Nagy, T. I., & Berend, G. (2011).** Multiword Expressions and Named Entities in the Wiki50 Corpus. *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, Association for Computational Linguistics, pp. 289–295.

Article received on 20/01/2018; accepted on 05/03/2018.
Corresponding author is Zuzana Nevěřilová.