Temporal Information Extraction with Cross-Language Projected Data

Przemysław Jarzębowski[1] and Adam Przepiórkowski[2,1]

[1] University of Warsaw
[2] Institute of Computer Science, Polish Academy of Sciences

p.jarzebowski@students.mimuw.edu.pl
adamp@ipipan.waw.pl

Abstract. This paper presents a method used for extracting temporal information from raw texts in Polish. The extracted information consists of the text fragments which describe events, the time expressions and the temporal relations between them. Together with temporal reasoning, it can be used in applications such as question answering or for text summarization and information extraction. First, a bilingual corpus was used to project temporal annotations from English to Polish. This data was further enhanced by manual correction and then used for inducing classifiers based on Conditional Random Fields (CRF) and a Support Vector Machine (SVM). For the evaluation of this task we propose a cross-language method that compares the system's results with results for different languages. It shows that the temporal relations classifier presented here outperforms the state of the art systems for English when using the macro-average F_1-measure, which is well suited for this multiclass classification task.

Keywords: temporal information, temporal relation, event extraction, word alignment

1 Introduction

One of the key elements of deep text understanding is the ability to process temporal information. Those parts of natural language texts which describe sequences of events often mention times of occurrence of these events. Being able to establish such a temporal relation between events and their occurrence times just by analysing a sentence would much enhance some NLP applications. Temporal reasoning, for example, is an essential part of many question answering systems. Information about events and their time of occurrence automatically extracted from sources such as news articles or Wikipedia would make it possible to answer a broad range of time-related questions. Furthermore, it would make it possible to infer relations between events. In text summarization, knowledge about the events mentioned might be a good indicator of the text's most significant or informative parts.

Depending on the target application of the extracted temporal information, the definition of an event can differ, and temporal relations can take a different set of values. TimeML [13] together with annotation guidelines created for TimeBank corpora [14] present a formalization of the temporal information extraction task. They specify which fragments of the text should be identified as events, time expressions and also defines types of temporal relations. This commonly used standard ([8,11,15]) is followed here.

The process of extracting temporal information can be split into three tasks: identification of time expressions, identification of events and classification of temporal relations between time expressions and events. Training supervised machine learning classifiers to solve the last two of them has proven to give the best results [17]. However, this method requires data with temporal annotation, which for Polish was not available. Manually creating annotation is an expensive and long process, so instead, a bilingual corpus and word alignment were used to project annotations from English to Polish. An example of a sentence with its temporal information projected from English to Polish is presented in Table 1.

The annotation for the English part of the bilingual corpus was created automatically by TIPSem [8] – a temporal information system for English. Next, the extracted events, time expressions and temporal relations between events and time expressions were projected to the Polish part of the corpus. The annotation obtained this way is noisy not only because of word alignment errors, but also because of misclassifications by TIPSem. Projection constraints were applied to limit both types of errors. The scarcity of errors in the projected events makes them acceptable to use as reference data for the further process, but the classified temporal relations contain relatively more errors. A part of them was manually corrected and the rest was used only to boost the classifier. Note that this work is only concerned with temporal relations between events and time expressions, and the annotation of such relations is much less time consuming than complete temporal annotation involving event annotation. The projection algorithm and the correction process are described in detail in Sec. 3.

Classifying the type of temporal relation between different events is a difficult task. When annotating the TimeBank corpus, the inter-annotator agreement on the type of temporal relation was $F_1 = 0.55$ [11]. The authors of the article argued that this low score was due to the large number of event pairs available for comparison, so it was difficult for annotators to spot all of the existing temporal relations. This low score has a great impact on the quality of the data used to build classifiers. In order to avoid this problem, the work presented here considers only the classification of temporal relations between events and temporal expressions, where the data is much more reliable, and does not consider temporal relations between events. Also, these temporal relations are much more significant for some of the applications mentioned above because one could use them to accurately put the events on a timeline.

The dataset thus created was used to induce two classifiers: an event classifier and a classifier of temporal relations. The first one uses Conditional Random Fields (CRF [7]), which is also used by TIPSem; the latter is based on Support

Vector Machine (SVM [3]). Details of the training of the classifiers and the text features used are described in Sec. 4. Unlike English, Polish does not enforce strict word order in sentences and is highly inflectional. This has a great impact on the features chosen for the classification. For extracting time expressions, a set of extraction rules was created and a rule-based shallow parsing system Spejd [4] was used. The defined extraction rules use both the lemma of a word and its morphosyntactic properties.

Section 5 presents the results of the evaluation of the event classifier and the temporal relation classifier. By projecting temporal relations across languages, a comparison is made between the results for the classification of temporal relations obtained here and the results of Evita [15] and TIPSem.

2 Related work

Application of the TimeML standard makes it possible to compare the results achieved here with those reported for state of the art systems, specifically the Evita system and the system that had the best score in the TempEval2 competition [17] in the events identification task – TIPSem. Evita integrates a rule-based approach with machine learning for event recognition and classification of temporal relations. TIPSem is based on machine learning, and the set of features it uses for classification is enriched with semantic roles. Just as in case of TIPSem, the current work follows the machine learning approach.

In the solution presented here, a word-aligned bilingual corpus was used to create a resource with temporal annotation in a new language, similarly to [16]. There, temporal annotation was projected from English to German to build classifiers for events, time expressions and temporal relations. On the other hand, in the work presented here, the projected data is used for inducing an event classifier, for boosting the classifier of temporal relations, but also to perform a cross-language comparison of systems. This comparison was based on the manual annotation of a small set of the temporal relations in the word-aligned bilingual corpus.

In comparison with TimeEval2 and its task of temporal relation classification, the work presented here focuses not only on assigning a temporal relation type but also on making a decision whether the temporal relation exists or not, as in [10]. In our work, for the purpose of evaluating the classified temporal relations, a macro-average F_1 measure is reported. Macro-average F_1 is the average of the F_1 scores computed for each of the types of temporal relation. Some of the types of temporal relations are much more frequent than others, and this measure ensures that the ability of the resulting classifier to classify all of them with high performance is included in the final score. To the best of our knowledge, this problem of the minority relation types was not addressed in previous work.

3 Creating temporal data

The annotation of events in the Polish part of the corpus was obtained by projecting annotations from the English part. For this purpose, word alignment between the two parts of a parallel corpus was performed. An example of the projection of temporal information using the word alignment from Fig. 1 is presented in Table 1.

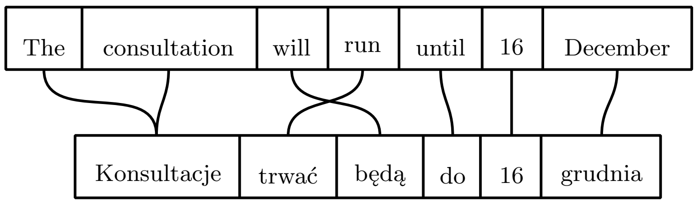

Fig. 1: Example of word alignment

Table 1: Example of temporal information projection using word alignment shown in Fig. 1

Sentence annotation:	*The* [*consultation*]$_{event_1}$ *will* [*run*]$_{event_2}$ *until* [*16 December*]$_{timex_1}$. [*Konsultacje*]$_{event_1}$ [*trwać*]$_{event_2}$ *będą do* [*16 grudnia*]$_{timex_1}$.
Temporal relations:	[*event₁*] *ENDED_BY* [*timex₁*] [*event₂*] *ENDED_BY* [*timex₁*]

3.1 Developing word alignment

Given one sentence written in two languages – a source sentence and its translation – word alignment, which looks at pairs of words across languages, finds those which have the corresponding meaning. It is not always a one-to-one relationship, and often for one English word multiple Polish words are found, and the other way around. The reasons for this are grammar differences between Polish and English (e.g. no determiners or phrasal verbs in Polish), and lexical differences (e.g. idioms).

Community Research and Development Information Service (CORDIS[3]) parallel corpus was used, and the word alignment was created for 146,334 of its sentences with a statistical machine translation tool – Moses [6]. Moses first builds a

[3] `http://cordis.europa.eu`

Hidden Markov Model for the entire corpus, and then for each sentence chooses the alignment with the highest probability, computing it with the built translation model. The final alignment is the result of merging two separate alignments: the Polish-English alignment and the English-Polish alignment. Both of them are of the type one-to-many. Merging them gives a many-to-many alignment reflecting the true relationship between words in those languages. Also, before the alignment was computed, the Polish part of the corpus was preprocessed, and all words were substituted for their lemmas using the tools Morfeusz [18] and Pantera [1]. The positive impact of this preprocessing step on the alignment accuracy was presented in [19]. Polish is a highly inflectional language, and the number of unique word types is much higher in the Polish part of the parallel corpus than it is in the English part. Because creating the word alignment is based on a simple string comparison, before the lemmatization all of the inflections of one Polish lemma were treated as different words. Lemmatization considerably increased the frequency of Polish-English word pairs, which helped in increasing the accuracy of the computed word alignment.

3.2 Projecting annotation

The developed word alignment was next used to project temporal information from English to Polish. `EVENT` and `TIMEX` tags, which cover the extent of events and the extent of time expressions, were copied alongside the word alignment. If the alignment was of the type one-to-many then the tag was multiplied. With the tags also their identifiers assigned by TIPSem were projected and, as a result, the temporal relations found by TIPSem became valid for the Polish part of the corpora. An example of a projection is shown in Table 1.

In order to limit the number of incorrect projections some constraints were applied. Some of them were suggested in [16]. Tag projections were required to be a contiguous sequence of words, and they were not allowed to clash, for example if two different events were projected to the same word. Unlike [16], the constraint on the `TIMEX` tags that they contain only content-bearing words (tokens which are not prepositions or punctuation) was not applied here. Prepositions, for example, are a valid part of a time expression, especially in the case of time expressions describing a duration, e.g. *from January to March 2011*. Sentences in which any of the constraints was not fulfilled were discarded from the dataset. Also, all the sentences which, after the projection, did not have any `EVENT` tag were omitted.

3.3 Annotating temporal relations

The temporal relations between events and time expressions were assigned one of the types: *before*, *ended by*, *after*, *begun by*, *is included*. *Is included* means that the event has happened in the time period defined by the given time expression. The TimeML annotation guidelines propose a total of 14 different temporal relations. However, some of them were symmetrical to the ones chosen here, and the other ones were not found useful for the aims of the presented work, i.e. finding temporal relations between events and time expressions. Also, TIPSem and Evita use the same types to describe temporal relations, but they also use the *simultaneous* type which is a special case of *is included*. During the annotation process, it was enforced that all the pairs of events and time expressions in a sentence were assigned one of the defined types, or *none* if there is no temporal relation. This addressed the problem of annotators accidentally omitting some of the temporal relations in the sentence. Those negative examples of temporal relations were also used when inducing the classifier, so that it is able to discriminate between the existence and non-existence of a temporal relation.

The types of temporal relations are unevenly distributed in text. The most frequent case is that there is no temporal relation between the given event and time expression. The relations *begun by*, *ended by* are rare (e.g. the TimeBank corpus for about 6.5K temporal relations has less than 1% relations *begun by*). In order to build a more balanced dataset, the preliminary projected temporal relations were used as a cue of their actual types. 187 sentences with 606 `TIMEX–EVENT` pairs were selected for manual annotation.

The events and time expressions used during the manual annotation are a result of the projection of temporal information found by TIPSem for English text, and they can contain some errors. A sentence with its `TIMEX–EVENT` pairs was discarded from the temporal relations corpora if the annotator discovered that an event which is in a temporal relation with the given time expression was not automatically detected. This way the annotators' work was limited to assigning a temporal relation rather than correcting the projected annotation and finding the actual events in the sentence. As a result, 240 temporal relation instances other than `none` were obtained, which is 4 times less examples than there were available for training in the TempEval2 contest for this task. That data was used both for training the temporal relations classifier and for evaluating the quality of data obtained with TIPSem.

4 Classification

4.1 Event recognition

In the literature two different approaches to the task of event recognition were proposed: one involved building a vocabulary of words describing events (e.g [2]), and the other approach used a machine learning classifier. One of the problems with the first approach is the property of language that a word which in one context describes an event, does not necessarily do so in another context. An example using the word *discover* is shown in the two sentences below:

- *The discovery of penicillin in 1928 by A. Fleming was a great breakthrough in medicine.*
- *The discovery of a cure for cancer would be a great breakthrough in medicine.*

For this reason a machine learning approach was applied which in comparison with using a set of defined rules can be much more resilient in such cases. The remaining part of this section introduces the method which is based on Conditional Random Fields.

Conditional Random Fields is a machine learning approach for sequence classification. Its main feature is that the classification process can use information about the class already assigned to the previous element in a sequence. In this application of CRFs the sequences are words arranged by their order of appearance in a sentence. Usually, two different events do not occur in a sentence as consecutive words, and if they do, they are a sequence of words with a specific relationship between them [9], as in *begin meeting*. This knowledge can be incorporated by CRFs classifier.

Text features. NLP for Polish is at a much less developed stage than it is for English, for example there is no robust syntactic parser available. For this reason, to represent the context features of words, the output of Spejd [4] – a rule-based shallow parsing tool – was used. Extraction rules were adopted to identify syntactic categories such as noun phrases and prepositional phrases. Also, morphosyntactic features such as the case of a word were used, because they carry some information about the word's semantic role in a sentence. For example the accusative case of a noun can mean that it is a patient of a verb. The features of words used for event recognition are:

- Lemma
- Polish WordNet [12] hypernyms of a lemma
- Grammatical features – part of speech (POS), case, gender, voice, tense, aspect, number
- Spejd features – syntactic category of a word, e.g. noun phrase, adjective phrase, adverb phrase, prepositional phrase
- Temporal expression proximity – often events are in close proximity to a time expression in a sentence. Those features contain information about whether the word is in the same dependent clause as a time expression, whether it occurs before or after one, and information about its distance from the time expression measured in number of words.

4.2 Temporal relations

The goal of the temporal classifier is to decide whether there is a temporal relation between the given time expression and event, and if there is one, then to classify its type. For this purpose, features of events and time expressions are used, as well as information about their relative position in a sentence. When developing the classifier the average of the individual F_1 values for all of the class types was maximised in a cross-validation process, so the resulting classifier can detect and classify a temporal relationship equally well regardless of its type.

SVM, which gives the best results for many machine learning tasks, was chosen here as a classification method. Each of the `TIMEX–EVENT` pairs was classified separately, and for that classification we did not find a strong use case for CRF classifier which can incorporate classification results of other data examples. Sampling was used when choosing the training data for each of the folds in order to guarantee that each of the data classes were of similar sizes.

Also a heuristic was implemented, which is applied if, for a given time expression, the classifier does not discover in the sentence any event with which that time expression is in a temporal relation. The heuristic is motivated by the assumption that at least one event in the sentence must be in a temporal relation with each time expression, i.e. dates in a sentence always describe the time of some event. Among the `TIMEX–EVENT` pairs with that time expression a pair which has the highest probability of some not *none* temporal relation type is found, and that temporal relation is assigned to that pair.

The following features were used for training the classifier of temporal relations:

- Features of time expressions – information about whether the time expression describes a specific date or a period of time, and prepositions preceding the time expression. Prepositions such as *after*, *while*, *until* often indicate the type of a temporal relation as is shown by experiments in [5].
- Features of events – morphological features of a head word describing the event, its tense and information about the syntactic category to which belongs.
- Features describing the relative position between events and time expressions – information about whether the word is in the same dependent clause as a time expression, whether it occurs before or after one, information about its distance from the time expression measured in number of words, and information about the syntactic categories on a path between the event and the time expression.

5 Evaluation

5.1 Event recognition

The evaluation of the event recognizer was conducted with 5-fold cross-validation on the projected data. Because this data was automatically obtained with TIPSem and word alignment, it is not free from misclassifications. These results are reported below in order to give an indication of how well the classifier is able to replicate TIPSem results. Table 2 compares the classifier's result with simple baselines which classify words as events using their POS and WordNet classes. Table 3 presents results broken down according to the POS of words denoting events.

Very low recall for adjectives can be explained by the fact that the dataset contains far less examples of events described with an adjective than with a noun or a verb. TIPSem performs worse in annotating those as well, which has an impact on the quality of the training and testing data.

Table 2: Event recognizer results against baselines

Classifier	Precision	Recall	F_1
Verbs only	51.3	70.3	59.3
Verbs + nouns selected with WordNet	42.1	82.7	55.8
Event recognizer	77.4	66.5	71.5

Table 3: Event recognizer results by POS

POS	Precision	Recall	F_1
Verb	79.9	83.1	81.5
Noun	62.4	32.6	42.8
Adjective	72.7	2.3	4.4
Other	75.0	7.1	13.0

5.2 Temporal relations

The evaluation of temporal relations was conducted with 5-fold cross-validation using manually annotated temporal relations (536 instances). Also, for the training of the classifier, a small sample of the automatically projected temporal relations was used in order to boost the classifier.

When reporting the performance of the classifier, the accuracy of the classifier's decision on whether the temporal relation exists or not is also considered. Although this was not assessed in the TempEval2 contest, we think that discriminating between these two situations is as important as deciding on the type of temporal relation. As well as the accuracy measure, the G_{mean} average and the average of the F_1 scores is reported for all the classes. G_{mean} is a geometrical mean of all the classes' recalls, and is frequently used to assess the quality of results when dealing with data unevenly distributed between classes. Maximising the average of the individual F_1 values guarantees that precision for all of the temporal relation types, as well as recall, will be included in the final score.

A baseline classifier following a simple algorithm was developed to compare the results obtained. The baseline for each time expression in a sentence chooses the event which is closest and assigns the temporal relation type based on the preposition before the time expression. If there is no preposition or the time expression has a duration type, then the *is included* relation is assigned. All other `TIMEX–EVENT` pairs with that time expression are assigned *none*. The results of the comparison are presented in Table 4.

The manually annotated temporal relations were also used to compare the performance of the classifier presented here with Evita and TIPSem. The events and time expressions used for annotation come from the projection of the temporal information found by TIPSem, so comparing its results with ours is straightforward. To compare the temporal relations found by Evita a mapping of its

Table 4: Results for the temporal relation classification

Classifier	*Accuracy*	G_{mean}	F_1avg
Baseline	74.7	43.8	53.4
Temporal relation wo heuristics	78.5	59.3	60.1
Temporal relation classifier	79.0	58.3	62.8

events and time expressions onto those found by TIPSem was applied. Some of the events were not recognised by Evita and vice versa, and as a result 276 events which matched the manually annotated temporal relations were found. By comparing the results across languages an assumption is made that the Polish translations in the parallel corpus does not change the type of temporal relation. This unfortunately is not always true, and some translations not following this rule were found. The comparison is presented in Table 5. The low G_{mean} and F_1avg score of TIPSem is due to its low performance for minority types of temporal relations, especially *after* and *before*.

Table 5: Comparison of the results for the classification of temporal relations across languages

Classifier	*Accuracy*	G_{mean}	F_1avg
TIPSem	73.7	0.0	24.4
Evita	66.7	35.2	44.6
Polish classifier	79.0	58.3	62.8

6 Conclusions and future work

This paper presents an approach to temporal information extraction from texts in languages which do not have dedicated corpora with temporal annotation. It uses the word alignment technique from the field of machine translation to create the required resources and then applies machine learning methods to train the event recognizer and the temporal relations classifier. The approach from [16] is extended here, and the manually annotated temporal relations are also used to directly compare performance of the presented system with the state of the art systems for English. This work shows that the effort to create resources in different languages for the task of temporal relations classification can benefit all of them. The manually annotated data in the Polish part of a parallel corpus can be projected to English and used as data for comparison of the systems.

In the presented work it is also proposed to maximise the macro-average F_1 measure when training a temporal relations classifier. This ensures that even the less frequent types of relation are classified with high performance.

The results of the evaluation show that the classification of temporal relations between `EVENT` – `TIMEX` pairs for Polish can be performed with relatively high accuracy just using prepositions. The applied machine learning approach which uses shallow parsing features of text improves that baseline, and significantly outperforms the temporal relations obtained by projecting data annotated with TIPSem and Evita. Those results suggest that either the task itself is easier for Polish language because of clearer relations between the preposition and the type of the temporal relation, or that the state of the art systems did not perform well in the classification of the minority types of temporal relations.

Future work should focus on enriching the corpus with manually annotated temporal relations, so that the less frequent types of temporal relations are represented by more examples. This could significantly increase the performance of their classification, as very few relations of those types were found automatically by the English systems.

References

1. Acedański, S.: A morphosyntactic Brill tagger for inflectional languages. In: Loftsson, H., Rögnvaldsson, E., Helgadóttir, S. (eds.) Advances in Natural Language Processing. Lecture Notes in Computer Science, vol. 6233, pp. 3–14. Springer (2010)
2. Arnulphy, B., Tannier, X., Vilnat, A.: Automatically generated noun lexicons for event extraction. In: Gelbukh, A. (ed.) Computational Linguistics and Intelligent Text Processing, Lecture Notes in Computer Science, vol. 7182, pp. 219–231. Springer Berlin / Heidelberg (2012)
3. Boser, B.E., Guyon, I.M., Vapnik, V.N.: A training algorithm for optimal margin classifiers. In: Proceedings of the Fifth Annual Workshop on Computational Learning Theory. pp. 144–152. ACM Press (1992)
4. Buczyński, A., Przepiórkowski, A.: Spejd: A shallow processing and morphological disambiguation tool. In: Vetulani, Z., Uszkoreit, H. (eds.) Human Language Technology. Challenges of the Information Society, Lecture Notes in Computer Science, vol. 5603, pp. 131–141. Springer Berlin / Heidelberg (2009)
5. Derczynski, L., Gaizauskas, R.: Using signals to improve automatic classification of temporal relations. In: Proceedings of the ESSLLI StuS (2010)
6. Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., Bojar, O., Constantin, A., Herbst, E.: Moses: Open source toolkit for statistical machine translation. In: Proceedings of the 45th Annual Meeting of the ACL on Interactive Poster and Demonstration Sessions. pp. 177–180. ACL '07, Association for Computational Linguistics, Stroudsburg, PA, USA (2007)
7. Lafferty, J.D., McCallum, A., Pereira, F.C.N.: Conditional Random Fields: Probabilistic models for segmenting and labeling sequence data. In: Proceedings of the Eighteenth International Conference on Machine Learning. pp. 282–289. ICML '01, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2001)

8. Llorens, H., Saquete, E., Navarro, B.: TIPSem (English and Spanish): Evaluating CRFs and Semantic Roles in TempEval-2. In: Proceedings of the 5th International Workshop on Semantic Evaluation. pp. 284–291. Association for Computational Linguistics, Uppsala, Sweden (2010)
9. Llorens, H., Saquete, E., Navarro-Colorado, B.: TimeML events recognition and classification: learning CRF models with semantic roles. In: Proceedings of the 23rd International Conference on Computational Linguistics. pp. 725–733. COLING '10, Association for Computational Linguistics, Stroudsburg, PA, USA (2010)
10. Llorens, H., Saquete, E., Navarro-Colorado, B.: Automatic system for identifying and categorizing temporal relations in natural language. International Journal of Intelligent Systems 27(7), 680–703 (2012)
11. Mani, I., Verhagen, M., Wellner, B., Lee, C.M., Pustejovsky, J.: Machine learning of temporal relations. In: Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics. pp. 753–760. ACL-44, Association for Computational Linguistics, Stroudsburg, PA, USA (2006)
12. Maziarz, M., Piasecki, M., Szpakowicz, S.: Approaching plWordNet 2.0. In: Proceedings of the 6th Global Wordnet Conference. Matsue, Japan (2012)
13. Pustejovsky, J., Castaño, J., Ingria, R., Saurí, R., Gaizauskas, R., Setzer, A., Katz, G.: TimeML: Robust specification of event and temporal expressions in text. In: Fifth International Workshop on Computational Semantics IWCS-5 (2003)
14. Pustejovsky, J., Hanks, P., Saurí, R., See, A., Gaizauskas, R., Setzer, A., Radev, D., Sundheim, B., Day, D., Ferro, L., Lazo, M.: The TimeBank corpus. In: Proceedings of Corpus Linguistics 2003. pp. 647–656 (2003)
15. Saurí, R., Knippen, R., Verhagen, M., Pustejovsky, J.: Evita: A Robust Event Recognizer for QA Systems. In: Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing. pp. 700–707. HLT '05, Association for Computational Linguistics, Stroudsburg, PA, USA (2005)
16. Spreyer, K., Frank, A.: Projection-based acquisition of a temporal labeller. In: Proceedings of IJCNLP 2008, Hyderabad, India. pp. 489–496 (2008)
17. Verhagen, M., Mani, I., Sauri, R., Knippen, R., Jang, S.B., Littman, J., Rumshisky, A., Phillips, J., Pustejovsky, J.: Automating temporal annotation with TARSQI. In: Proceedings of the ACL 2005 on Interactive poster and demonstration sessions. pp. 81–84. ACLdemo '05, Association for Computational Linguistics, Stroudsburg, PA, USA (2005)
18. Woliński, M.: Morfeusz — a Practical Tool for the Morphological Analysis of Polish. In: Klopotek, M., Wierzchon, S., Trojanowski, K. (eds.) Intelligent Information Processing and Web Mining, Advances in Soft Computing, vol. 35, pp. 511–520. Springer Berlin / Heidelberg (2006), http://dx.doi.org/10.1007/3-540-33521-8_55
19. Wróblewska, A.: Polish-English Word Alignment: Preliminary Study. In: Ryżko, D., Rybiński, H., Gawrysiak, P., Kryszkiewicz, M. (eds.) Emerging Intelligent Technologies in Industry, Studies in Computational Intelligence, vol. 369, pp. 123–132. Springer Berlin / Heidelberg (2011)