

Named Entity Recognition in Islam-related Russian Tweets

Valeria Mozharova, Natalia Loukachevitch

Lomonosov Moscow State University,
Russia

joinmek@rambler.ru, louk.nat@mail.ru

Abstract. The paper describes an approach to creating a domain-specific tweet collection written by users frequently discussing Islam-related issues in Russian. We use this collection to study specific features of named entity recognition on Twitter. We found that in contrast to tweets collected randomly, our tweet collection contains relatively small number of spelling errors or strange word shortenings. Specific difficulties of our collection for named entity recognition (NER) include a large number of Arabic and other Eastern names and frequent use of ALL-CAPS spelling for emphasizing main words in messages. We studied the transfer of the NER model trained on a newswire collection to the created tweet collection and approaches to decrease the degradation of the model because of the transfer. We found that for our specialized text collection, the most improvement was based on normalized word capitalization. Two-stage approaches to named entity recognition and word2vec-based clustering were also useful for our task.

Keywords: Islam russian, NER model, russian tweets, normalization of tweets.

1 Introduction

Available electronic texts in social media allows the study of the differences in language or style for specific communities [10, 16], or the dependence of the language on social and demographic characteristics of users [11, 23].

In Russia, one of the prominent groups of population are Muslims; the number of Islam believers is estimated as 10-30% of the total population. To study the language and the thematic interests of Muslims in Twitter, we created a specialized tweet collection. We begin our work with this collection from named entity recognition (NER) on Islam-related tweets.

Most studies of NER have been carried out on news collections and shown high quality of named entity extraction. However, the transfer of NER recognizers to other genres of texts usually demonstrates a significant decrease in the performance.

In particular, application of general NER recognizers designed for or trained on news collections can show the decrease in performance of up to 50% or more on informal texts published on social media platforms such as Twitter or Facebook [3, 6, 7, 8, 20].

In this paper we consider the transfer of a Russian NER recognizer trained on news texts to extracting names from Twitter messages. Our tweet collection is specialized, it is gathered from messages of those users who discuss issues related to Islam in their posts in contrast to other studies where Twitter collections are formed with random sampling of Twitter messages.

This allows us to reveal specific features of the tweet language of Islam-oriented and other similar communities. We consider the transfer of a CRF-based NER recognizer from a news data to the tweet collection and approaches to decrease the degradation of the model due to the transfer.

2 Related Work

2.1 Named Entity Recognition for Twitter

It is known that extraction of names from Twitter messages is a much more difficult task than from other genres of texts because of their shortness and informal character. In [6] the authors review the problems and approaches to named entity recognition and entity linking for tweets. They write that the tweet content is noisy because of incorrect spelling, irregular capitalization, and unusual abbreviations. In their experiments, the main sources of mistakes in named entity recognition in tweets were violations in capitalization, especially a large number of names written in lower case. They studied automatic normalization of tweets including spelling and capitalization correction and reported that in their investigation the normalization slightly improved the performance in NER for tweets.

In [19] the authors write that due to unreliable capitalization in tweets, common nouns are often misclassified as proper nouns, and vice versa. Some tweets contain all lowercase words (8%), whereas others are in ALL CAPS (0.6%). In addition to differences in vocabulary, the grammar of tweets differs from news text, for example, tweets often start with a verb. In their experiments the supervised approach was used to predict correct capitalization of words. The set of features included: the fraction of words in the tweet which are capitalized, the fraction which appear in a dictionary of frequently lowercase/capitalized words but are not lowercase/capitalized in the tweet, the number of times the word 'I' appears lowercase and whether or not the first word in the tweet is capitalized.

To study NER on Twitter performed with several NER systems, the authors of [7] use crowdsourcing to annotate tweets for the NER task. They annotate all @usernames as PER (person name). Annotating tweets for their experiments, the authors of [19] chose not to annotate @usernames mentioned in tweets as entities because it is trivial to identify them using a simple regular expression, and they would only serve to inflate the performance statistics.

In [3] the authors study the transfer of their NER model from news texts to tweets. They create a training set consisting of 1000 tweets. They use a baseline NER model based on token and context features (wordform, lemma, capitalization, prefixes and suffixes) and enhance it with two unsupervised representations (Brown clusters and vector representations) based on a large collection of unannotated tweets.

Besides, they propose a technique to combine a relatively small Twitter training set and larger newswire training data. They report that two unsupervised representations work together better than alone, and the combination of training sets further improves the performance of their NER system.

2.2 Named Entity Recognition in Russian

In Russian there is a tradition of engineering approaches to the named entity recognition task [12, 13, 22]. Machine-learning approaches for Russian NER usually employ the CRF machine learning method.

In [1] the authors presented the results of the CRF-based method on several tasks, including named entity recognition. The experiments were carried out on their own Russian text corpus, which contained 71,000 sentences. They used only n-grams and orthographic features of tokens without utilizing any knowledge-based features. They achieved 89.89% of F-score on three named entity types: persons (93.15%), geographical objects (92.7%), and organizations (83.83%).

In [17] the experiments utilized the open Russian text collection "Persons-600"¹ for the person name recognition task. The CRF-based classifier employed such features as token features, context features, and the features based on knowledge about persons (roles, professions, posts, and other). They achieved 88.32% of F-score on person names.

In [9] the experiments were carried out on the Russian text collection, which contained 97 documents. The authors used two approaches for the named entity recognition: knowledge-based and CRF-based approach. In the machine learning framework they utilized such features as the token features and the knowledge features based on word clustering (LDA topics [4], Brown clusters [2], Clark clusters [5]). They achieved 75.05% of F-score on two named entity types: persons (84.84%) and organizations (71.31%).

In 2016 the FactRuEval competition for the Russian language was organized. The FactRuEval tasks included recognition of names in Russian news texts, recognition of specific attributes of names (family name, first name, etc), and extraction of several types of facts [21]. The named entity recognition task on tweets has not been studied before for Russian. Also, the dependence of NER performance on the language of specific Twitter user communities has not been researched yet.

3 Text Collections

3.1 News Text Collection

We study the transfer of CRF-based NER classifier trained on newswire data to the tweet collection. For training our system, we chose open Russian text collection "Persons-1000", which contains 1000 news documents labeled with three types of named entities: persons, organizations and locations².

¹ ai-center.botik.ru/Airec/index.php?option=com_content&view=article&id=27:persons-600&catid=15&Itemid=40

² labinform.ru/pub/named_entities/descr_ne.htm

Table 1. The quantitative characteristics of the labeled named entities in text collections.

Type	News collection	Twitter collection
PER	10623	1546
ORG	8541	1144
LOC	7244	2836
OVERALL	26408	5526

The labeling rules are detailed in [15]. The counts of each named entity type in the collection are listed in Table 1.

3.2 Tweet Text Collection

We are interested in study of the language of Islam-related Twitter users in Russian. To extract tweets from users discussing Islam-related issues, we created a list of 2700 Islam terms. Then we extracted Russian tweets mentioning these terms using Search Twitter API, got users' accounts containing extracted tweets and ordered the accounts in the decreased number of extracted tweets from these accounts. We found that a lot of words from our list practically are not mentioned in tweets, other words (for example, "mosque" or "muslim") are often used by very different people, not only Muslims.

After studying tweets from the extracted accounts, we created a very small list of the main Islam words ("Allah", "Quran", "Prophet", in various forms of Russian morphology). We also added the names of several known islamist organizations to find their possible non-muslim proponents. Then we repeated the whole procedure of tweet and account extraction, and found that the extracted collections can be considered as an appropriate approximation of messages generated by Islam-related users.

We selected 100 users with the largest number of the extracted tweets, downloaded all their tweets and obtained the tweet collection consisting of 300 thousand tweets (further *FullTweetCollection*). Then we randomly extracted tweets from different users, removed non-Russian or senseless tweets and at last obtained the tweet collection of 4192 tweets (further *TestTweetCollection*).

The created collection contains messages with Quran quotes, religious and political argumentation, news-related messages mainly about Near and Middle East events (Syria, Iraq, Afghanistan etc) and Islamist organizations (Syrian opposition groups, ISIL, etc.) and also other types of messages (for example, advertisements).

The obtained collection was labeled similar to "Persons-1000". To annotate numerous mentions of Allah, we added the Deity type to the annotation scheme, but in the current study we consider the Deity type as a subtype of the Person type.

Analyzing the created collection from the point of view of NER difficulties we found that violations in capitalization mainly include ALL-CAPS words for the whole tweet and its fragment. Such capitalization is used for emphasizing important words in the text or words related to Allah as in the following example: За все потери, ОН дает нам большую награду" ("For all the losses He gives us a great reward").

Also the tweets mention a lot of Eastern names of persons, organizations ("Фастаким кама умирт"(Fastakim Kama Umirt group), Джабхат фатх аш-Шам (Jabhat Fateh al-Sham)), or local places difficult for correct recognition.

The fraction of tweets with spelling mistakes, unusual shortenings is relatively low. We suppose that this is because the selected users are well-educated, they are professional writers in some sense, in most cases they are leaders of opinions, whose messages are retweeted by many other people. Therefore it is especially useful to study the specific features of their tweet language.

4 Description of NER Model

In our study we employ the baseline CRF-classifier that utilizes token features, context features, and lexicon features for NER. Then we consider the ways to improve the baseline model adapting it to the Twitter language. The adaptation techniques include the use of two stage-processing and unsupervised word clustering. Besides, we test the impact of tweet normalization on the NER performance.

4.1 Baseline Model

Before named entity recognition with CRF, tweets are processed with a morphological analyzer for determining the part of speech, gender, lemma, grammatical number, case and other characteristics of words. This information is used to form features of each word for classifying. In the baseline model we consider three types of features: local token features, context features, and features based on lexicons.

Token Features: The token features include:

- Token initial form (lemma),
- Number of symbols in a token,
- Letter case. If a token begins with a capital letter, and other letters are small then the value of this feature is "BigSmall". If all letters are capital then the value is "BigBig". If all letters are small then the value is "SmallSmall". In other cases the value is "Fence",
- Token type. The value of this feature for lexemes is the part of speech, for punctuation marks the value is the type of punctuation,
- Symbol ngrams. The presence of prefixes, suffixes and other ngrams from the predefined sets in a token.

Context-based Features The group of context features includes two feature types. The first type is local context features. It takes into account all mentioned token feature values of nearby words in two-word window to the right and to the left from the current word.

The second type is the bigram context feature. It contains information about the determined named entity type of the previous word. It helps to find named entity borders more precisely. For example, if the person second name is difficult for recognition, the presence of the first name before this word makes the classification easier.

Table 2. Vocabulary sizes.

Vocabulary	Size, objects	Clarification
Famous persons	31482	Famous people
First names	2773	First names
Surnames	66108	Surnames
Person roles	9935	Roles, posts
Verbs of informing	1729	Verbs that usually occur with persons
Companies	33380	
Company types	6774	Organization names
Media	3909	Organization types
Geography	8969	Media
Geographical adjectives	1739	Geographical objects
Usual words	58432	Geographical adjectives
Equipment	44094	
		Frequent Russian words (nouns, verbs, adjectives)
		Devices, equipment, tools

Features based on Lexicons: To improve the quality of recognition, we utilize special lexicons with lists of useful objects. An object can be a word or a phrase. The lexicons had been created before the current work and were not changed during the study.

To calculate the lexicon features, the system matches the text and lexicon entries. If a token is met in a matched lexicon entry then it obtains the lexicon feature value equal to the word length of the found entry. The use of the entry length as a feature helps to diminish the affect of lexical ambiguity.

For example, in the list of organizations there is “Apple” as the name of a company. But this word does not necessarily mean a company because it has the second “fruit” sense. In the opposite, if the phrase “Lomonosov Moscow State University” is found in the text, the probability of the organization sense is higher than in the first case. The lexicon feature containing the matched entry length helps the system to distinguish these two cases.

The biggest lexicons are listed in Table 2. The overall size of all vocabularies is more than 335 thousand entities. These lexicons were collected from several sources: phonebooks, Russian Wikipedia, RuThes thesaurus of the Russian language [14].

4.2 Adaptation of NER Model to Tweets

Unsupervised Word Clustering: In previous studies it was shown that unsupervised word clustering on the basis of a large text collection improves the NER performance [18, 15]. In our case we compare the impact of word clusters calculated on a large news collection and large tweet collection. For clustering we use the word2vec package¹. It represents words with distributional vectors (word embeddings), computed with the neural network.

¹ github.com/dav/word2vec

The semantic similarity between two words is calculated as the cosine measure between two corresponding vectors. The package allows automatic clustering of words according to their calculated similarity. We used the cbow model with vector sizes equal to 300. Thus, each word has an additional feature – the number of a cluster in that it appears. The news collection utilized for clustering contains two million news documents. For tweet-based clustering we use a tweet collection consisting of randomly extracted Russian tweets and including 8.3 million tweets.

Two-stage Prediction: We suppose that for adapting a classifier to a text collection it can be useful to take into account the entities already labeled by the classifier and to memorize the named entity type statistics for future use.

At the first stage, the classifier extracts named entities. Then the system collects the class statistics determined at this stage for each word and uses it to form features of the second stage. After that, new features together with old ones participate in final classification. These statistics can be collected from the current text (the whole text or its part preceding to the word under analysis) or from a set of texts (collection statistics). In case of tweet processing, texts are small therefore only the collection statistics can be used. In our experiments this statistics can be obtained from the FullTweetCollection gathered from the selected user accounts or the labeled TestTweetCollection as described in Section 3.2.

For each word, the system finds all mentions of this word in the processed collection and counts frequencies of determined named entity types for this word. Using these frequencies for each entity type, the system creates additional features, which have one of three values: *no_one* (if the word has not been recognized as a named entity of the chosen type), *best* (if the word has been assigned to the chosen named entity type more than in 50% of cases), and *rare* (if the word has been assigned to the chosen named entity type less than in 50% of cases). For example, if the word "Russia" was met 500 times in a collection, and the classifier assigned it 200 times to organizations and 300 times to locations, then the values of the global statistics feature for the word "Russia" will be as following: *PER* – *no_one*, *ORG* – *rare*, *LOC* – *best*.

4.3 Normalization of Word Capitalization

As we found that in our tweet collection the share of misprints is not very high we did normalization only for word capitalization. The normalization was based on the large news collection described in Section 4.2. For each word in this collection, we counted how many times the word was written in letter case or capital case when it stands not in the beginning of a sentence. The more frequent case was considered as normal for this word. We considered the normalization in two variants:

- Variant A. All words in a tweet, except the first one, are changed to a normal form of capitalization,
- Variant B. All words in a tweet including the first one are changed to a normal form of capitalization.

We found that the variant B produces better results and later experimented only with this variant.

#ДЕУНАК РЕЛИГИИ НА ЗЕМЛЕ ТОЖЕ СТАРЕЮТ\СЕЙЧАС ОСНОВНАЯ ЭТО КАТОЛИКИ
НАБИРАЕТ СИЛУ И СТАНЕТ ЕДИНОЙ ВСЕОБЩЕЙ ИСЛАМ\ОСТАЛЬНЫЕ
ОТМИРАЮТ\ПРОРОК

#ДЕУНАК США КОНСЕРВАТОРЫ ХОТЯТ ПОВЕРНУТЬ ВРЕМЯ НАЗАД\КИТАЙ ПОДЛЕЖИТ
ЗАМЕНЕ ИЛИ УНИЧТОЖЕНИЮ\ИНДИЯ ИСЛАМ\РУССКИЕ ВЫМРУТ БУДЕТ
ИСЛАМ\ПРОРОК ДЕ

#ДЕУНАК ЧЕЛОВЕЧЕСТВО БЕЗ МОЕГО УЧАСТИЯ МОЖЕТ ОТСРОЧИТЬ ВСЕМИРНЫЙ
ПОТОП СОКРАТИВ НА ЗЕМЛЕ КОЛИЧЕСТВО РУССКИХ ДО 50 МИЛЛИОНОВ\ПРОРОК
ДЕЙНАК!!

Fig. 1. Tweets before normalization.

ДЕУНАК религии на земле тоже стареют\сейчас основная это католики
набирает силу и станет ЕДИНОЙ всеобщей ислам\остальные отмирают\пророк

ДЕУНАК США консерваторы хотят повернуть время назад\КИТАЙ подлежит
замене или уничтожению\индия ислам\русские вымрут будет ислам\пророк де

ДЕУНАК человечество без моего участия может отсрочить ВСЕМИРНЫЙ потоп
сократив на земле количество русских до 50 миллионов\ПРОРОК ДЕЙНАК!!

Fig. 2. Tweets after normalization.

Fig. 1 presents several tweets with the manual annotation before normalization. Fig. 2 shows the same tweets after normalization. Also the hashtag symbols were removed from a word if this word was found in the news collection to improve its matching with the lexicons.

5 Experiments

In preprocessing we removed mentioned user accounts with “@” and urls in the end of tweets. We consider these data as additional, as meta-information, from which we should not extract names. We train the described variants of our NER model on the news collection. Table 3 shows the results of named entity recognition on the “Persons-1000” collection (cross-validation 3:4).

It can be seen that our baseline model is quite good on the news collection and slightly improved after adding clustering features and the two-step approach. In this case the collection statistics is obtained from the same “Persons-1000” collection. Then we apply the trained model to the test tweet collection in initial capitalization and normalized capitalization.

Table 3. News collection NER performance.

Model	F-measure, %
Baseline	92.49
Baseline + News clusters	93.48
Baseline + News clusters + Collection statistics	93.53

Table 4. Tweet NER performance.

Model	F-measure, TestTweetCollection	F-measure, Normalized TestTweet-Collection
1) Baseline	64.44%	69.88%
2) Baseline + Collection statistics (TestTweetCollection)	64.99%	70.32%
3) Baseline + Collection statistics (FullTweetCollection)	65.78%	70.44%
4) Baseline + news clusters	66.03%	70.88%
5) Baseline + tweet clusters	66.08%	70.36%
6) Baseline + tweet and news clusters	66.23%	70.89%
7) (2) + tweet and news clusters	67.27%	71.20%
8) (3) + tweet and news clusters	66.46%	69.73%

Table 4 presents the performance of NER models trained on the "Persons-1000" collection for the tweet data. One can see that all models significantly degrade on the tweet collection.

The normalization significantly improves the performance of NER (in contrast to other studies [6]). Word clustering and the collection statistics improve both NER for initial and normalized text collections. Their impact is larger than for the news collection (Table 3).

The combination of tweet and news clusters was better than only tweet clusters possibly because of the political and religious character of the gathered collection. In total, the NER performance improves more than 10% on the tweet data.

Analyzing the results generated with the best model on the normalized collection we can see still significant share of mistakes because of incorrectly normalized capitalization. The main subtypes of such problems include:

- Ambiguous words with different capitalization ("Earth", "Rose") in different contexts.
- Words that should be capitalized in this specific collection. For example, "Paradise" and "Hell" seem to be specific entities in this genre of texts.
- Multiword expressions in that each word is usually written in letter case, but the whole multiword expression denotes a name and at least the first its word should be capitalized.

For example, the expression "Московская область" (*Moscow region*) is normalized incorrectly because the word "московский" is written in letter case more frequently in the Russian news collection. However, "Московская область" is the name of the location entity and should be capitalized.

6 Conclusion

The paper describes an approach to creating a domain-specific tweet collection written by users frequently discussing Islam-related issues in Russian. We use this collection to study specific features of named entity recognition on Twitter.

We found that in contrast to tweets collected randomly, our tweet collection contains relatively small number of spelling errors or strange word shortenings. Specific difficulties of our collection for named entity recognition include a large number of Arabic and other Eastern names (persons, locations, organizations) and frequent use of ALL-CAPS writing for emphasizing main words in messages.

We have studied the transfer of NER model trained on a newswire collection to the created tweet collection and approaches to decrease the degradation of the model because of the transfer. We found that for our specialized text collection, the most improvement was based on normalizing of word capitalization. Two-stage approaches to named entity recognition and word2vec-based clustering were also useful for our task. In future we plan to improve techniques of tweet normalization and study NER for tweets of followers of the selected users.

Acknowledgment. This work is partially supported by RFBR grant No. 16-29-09606.

References

1. Antonova, A., Soloviev, A.: Conditional random field models for the processing of russian. In: Computational Linguistics and Intellectual Technologies: Papers From the Annual Conference "Dialogue", vol. 1, pp. 27–44 (2013)
2. Brown, P. F., Desouza, P. V., Mercer, R. L., Della-Pietra, V. J., Lai, J. C.: Class-based n-gram models of natural language. *Computational Linguistics*, vol. 18, no. 4, pp. 467–479 (1992)
3. Cherry, C., Guo, H.: The unreasonable effectiveness of word representations for Twitter named entity recognition. In: Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 735–745 (2015) doi: 10.3115/v1/N15-1075
4. Chrupala, G.: Efficient induction of probabilistic word classes with LDA. In: Proceedings of the 5th International Joint Conference on Natural Language Processing, pp. 363–372 (2011)
5. Clark, A.: Combining distributional and morphological information for part of speech induction. In: 10th Conference of the European Chapter of the Association for Computational Linguistics, vol. 1, pp. 59–66 (2003) doi: 10.3115/1067807.1067817
6. Derczynski, L., Maynard, D., Rizzo, G., Van-Erp, M., Gorrell, G., Troncy, R., Petrak, J., Bontcheva, K.: Analysis of named entity recognition and linking for tweets. *Information Processing and Management*, Elsevier, vol. 51, no. 2, pp. 32–49 (2015) doi: 10.1016/j.ipm.2014.10.006

7. Finin, T., Murnane, W., Karandikar, A., Keller, N., Martineau, J., Dredze, M.: Annotating named entities in Twitter data with crowdsourcing. In: *Proceedings of the NAACL Workshop on Creating Speech and Text Language Data With Amazon's Mechanical Turk*, pp. 80–88 (2010)
8. Fromreide, H., Hovy, D., Søgaaard, A.: Crowdsourcing and annotating NER for Twitter #drift. In: *Proceedings of the Ninth International Conference on Language Resources and Evaluation*, pp. 2544–2547 (2014)
9. Gareev, R., Tkachenko, M., Solovyev, V., Simanovsky, A., Ivanov, V.: Introducing baselines for Russian named entity recognition. In: *Lecture Notes in Computer Science*, vol. 7816, pp. 329–342 (2013) doi: 10.1007/978-3-642-37247-6_27
10. Hidayatullah, A. F.: Language tweet characteristics of indonesian citizens. In: *International Conference on Science and Technology (TICST)*, pp. 397–401 (2015) doi: 10.1109/TICST.2015.7369393
11. Hovy, D.: Demographic factors improve classification performance. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, vol. 1, pp. 752–762 (2015) doi: 10.3115/v1/P15-1073
12. Khoroshevsky, V. F.: Ontology driven multilingual information extraction and intelligent analytics. *NATO Science for Peace and Security Series*, vol. 27, pp. 239–262 (2010) doi: 10.3233/978-1-60750-611-9-239
13. Kuznetsov, I. P., Kozerenko, E. B., Kuznetsov, K. I., Timonina, N. O.: Intelligent system for entities extraction (ISEE) from natural language texts. In: *Proceedings of the International Workshop on Conceptual Structures for Extracting Natural Language Semantics-Sense*, vol. 9, pp. 17–25 (2009)
14. Loukachevitch, N., Dobrov, B. V.: RuThes linguistic ontology vs. russian wordnets. In: *Proceedings of the seventh global wordnet conference*, pp. 154–162 (2014)
15. Mozharova, V. A., Loukachevitch, N. V.: Combining knowledge and CRF-based approach to named entity recognition in russian. In: *International Conference on Analysis of Images, Social Networks and Texts, Communications in Computer and Information Science*, pp. 185–195 (2016) doi: 10.1007/978-3-319-52920-2_18
16. Paris, C., Thomas, P., Wan, S.: Differences in language and style between two social media communities. In: *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*, vol. 6 (2012) doi: 10.1609/icwsm.v6i1.14307
17. Podobryaev, A.: Person names recognition using CRF model. In: *15th All-Russian Scientific Conference Digital Libraries: Advanced Methods and Technologies, Digital Collections*, pp. 255–258 (2013)
18. Ratinov, L., Roth, D.: Design challenges and misconceptions in named entity recognition. In: *Proceedings of the Thirteenth Conference on Computational Natural Language Learning*, pp. 147–155 (2009)
19. Ritter, A., Clark, S., Etzioni, O., Mausam: Named entity recognition in tweets: an experimental study. In: *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pp. 1524–1534 (2011)
20. Ritter, A., Etzioni, O., Clark, S., Mausam: Open domain event extraction from Twitter. In: *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1104–1112 (2012) doi: 10.1145/2339530.2339704
21. Starostin, A., Bocharov, V., Alexeeva, S., Bodrova, A., Chuchunkov, A., Dzhumaev, S., Efimenko, I., Granovsky, D., Khoroshevsky, V., Krylova, I., Nikolaeva, M., Smurov, I., Toldova, S.: FactRuEval 2016: Evaluation of named entity recognition and fact extraction systems for Russian. In: *Proceedings of the International Conference "Dialogue 2016", Computational Linguistics and Intellectual Technologies* (2016)

Valeria Mozharova, Natalia Loukachevitch

22. Trofimov, I.: Person name recognition in news articles based on the persons-1000/1111-f collections. In: 16th All-Russian Scientific Conference Digital Libraries: Advanced Methods and Technologies, Digital Collections, pp. 217–221 (2014)
23. Yang, Y., Eisenstein, J.: Putting things in context: Community-specific embedding projections for sentiment analysis. Arxiv-Social Media Intelligence (2015)