

Distant Supervision for Sentiment Attitude Extraction

Nicolay Rusnachenko¹, Natalia Loukachevitch^{1,2}, Elena Tutubalina^{3,4}

¹ Bauman Moscow State Technical University, Moscow, Russia

² Lomonosov Moscow State University, Moscow, Russia

³ Kazan Federal University, Kazan, Russia

⁴ Samsung-PDMI Joint AI Center, PDMI RAS, St. Petersburg, Russia

kolyarus@yandex.ru, louk_nat@mail.ru, elvtutubalina@kpfu.ru

Abstract

News articles often convey attitudes between the mentioned subjects, which is essential for understanding the described situation. In this paper, we describe a new approach to distant supervision for extracting sentiment attitudes between named entities mentioned in texts. Two factors (pair-based and frame-based) were used to automatically label an extensive news collection, dubbed as RuAttitudes. The latter became a basis for adaptation and training convolutional architectures, including piecewise max pooling and full use of information across different sentences. The results show that models, trained with RuAttitudes, outperform ones that were trained with only supervised learning approach and achieve 13.4% increase in F1-score on RuSentRel collection.¹

1 Introduction

Relation extraction nowadays remains one of the popular tasks in the natural language processing domain. The relation types to be extracted from texts may vary and result in different tasks: semantic classification of relations between a pair of common nominals (Hendrickx et al., 2009), source-target sentiment relation extraction (Ellis et al., 2014), opinion expression towards entities and events (Deng and Wiebe, 2015a), attitude extraction between mentioned named entities (Rusnachenko and Loukachevitch, 2018), etc.

Dealing with one of these tasks, the greatest difficulty one encounters is the complexity of the sentence structure. As for analytical articles, the idea expressed by the author could be conveyed in different variants, which is a feature of natural languages.

¹The code is available on <https://github.com/nicolay-r/attitudes-extraction-ds>

When relation extraction is performed automatically using machine learning approaches, this complexity results in a lack of training examples. One technique that helps to accomplish this task is *distant supervision* (DS), initially proposed by (Mintz et al., 2009). It assumes to extract and label data by relying on assumptions based on a prepared knowledge base. Although many methods have been proposed in such domains as sentiment analysis and relation extraction (Turney, 2002; Zeng et al., 2015), the domain of sentiment attitude extraction remains understudied.

This paper describes a new approach to distant supervision for extracting sentiment attitudes between named entities mentioned in texts. It is worth noting that DS faces the problem of wrong labels, which becomes a reason of noisy labeled data. To address the shortcomings of noisy labeling, in this paper we exploit two primary sources of automatic annotation:

- Prior knowledge about current attitudes between political entities (figures);
- Sentiment frames that define attitudes between participants of a situation.

The obtained corpus annotated with attitudes was used to train convolutional neural networks (CNNs), adapted for relation extraction and full use of information across multiple sentences. Our key contributions in this work are two-fold:

- We propose a workflow of automatic sentiment attitudes extraction, which exploits news title simplicity to perform annotation;
- We conduct extensive experiments on RuSentRel (Loukachevitch and Rusnachenko, 2018) and the results demonstrate that CNNs trained on two types of training data achieve F1-score increase by 13.4% over models that do not employ obtained corpus in training.

2 Related Work

The task of attitude recognition toward named entities or events, including opinion holder identification from full texts did not attract much attention. In 2014, the TAC evaluation conference in Knowledge Base Population (KBP) track included so-called sentiment track (Ellis et al., 2014). The task was to find all the cases where a query entity (sentiment holder) holds a positive or negative sentiment about another entity (sentiment target). Thus, this task was formulated as a query-based retrieval of entity-sentiment from relevant documents and focused only on query entities².

MPQA 3.0 (Deng and Wiebe, 2015b) is a corpus of analytical articles with annotated opinion expressions (towards entities and events). The annotation is sentence-based. For example, in the sentence «When the Imam issued the fatwa against Salman Rushdie for insulting the Prophet ...», Imam is negative to Salman Rushdie but is positive to the Prophet.

In paper (Choi et al., 2016), authors studied the approach to the recovery of the documents attitudes between subjects mentioned in the text. The approach considers such features as frequency of a named entity in the text, relatedness between entities, direct-indirect speech, etc. The best quality of opinion extraction obtained in this work was only about 36% F-measure, which shows that the necessity of improving extraction of attitudes at the document level is significant and this problem has not been sufficiently studied.

A corpus of analytical articles, obtained from authoritative foreign sources and translated into Russian has been invented in (Loukachevitch and Rusnachenko, 2018). The collected articles contain both the author’s opinion on the subject matter of the article and a large number of attitudes mentioned between the participants of the described situations. Authors experiment with automatic attitudes extraction within the developed corpus. In comparison with (Choi et al., 2016) where documents much smaller and written in English, authors mentioned the closest F-measure and conclude that the task still remains complicated.

Each attitude may be considered in terms of related article context, or sentence. The sentence consists of words which could be gathered and treated as an *embedding*, where each word repre-

²<https://tac.nist.gov/2014/KBP/Sentiment/index.html>

sents a feature vector. Convolving embedded sentence representation by a set of different filters, in paper (Zeng et al., 2014) authors implemented and trained the Convolutional Neural Network (CNN) model for the relation classification task. Being applied for the SemEval-2010 Task 8 (Hendrickx et al., 2009), the obtained model significantly outperformed the results of other participants.

This idea was developed further in terms of *max-pooling* operation (Zeng et al., 2015). This is an operation, which is applied to the convolved by filters data and extracts the maximal values within each convolution. However, for the relation classification task, original max-pooling reduces information extremely rapid and blurs significant relation aspects. Authors proposed to treat each convolution in parts. The division into parts depends on attitude entities: *inner* (between entities), and *outer*. This approach resulted in an advanced architecture model and was dubbed as “Piecewise Convolutional Neural Network” (PCNN).

3 Resources

This section describes resources (collections and lexicons) that were used for the dataset annotation.

3.1 RuSentRel Collection

In our experiments, we use the RuSentRel corpus³ consisted of analytical articles from Internet-portal inosmi.ru (Loukachevitch and Rusnachenko, 2018) devoted to international relations. In this corpus, the manual annotation of the sentiment attitudes towards mentioned named entities had been carried out at the document level. The annotation is subdivided into two subtypes:

- The author’s relation to mentioned named entities;
- The relation of subjects expressed as named entities to other named entities.

An analytical document can refer to an entity with several variants of naming (*Vladimir Putin – Putin*), synonyms (*Russia – Russian Federation*), or lemma variants generated from different word-forms. For correct inference of attitudes between named entities in the whole document, the corpus is provided with a list of variant names for the same entity found in the corpus.

³<https://github.com/nicolay-r/RuSentRel/tree/v1.1>

In this paper, we utilize RuSentRel corpus in experiments for the proposed approach. Table 1 describes the corpus statistics.

Parameter	Value
Number of documents	73
Total opinion pairs	1361
Sentences (avg./doc.)	105.75
Opinion pairs (avg./doc.)	18.64
Positive opinion pairs (avg./doc.)	8.71
Negative opinion pairs (avg./doc.)	9.93
Avg. dist. between named entities within a sentence in words	10.2

Table 1: Attitude statistics of RuSentRel-v1.1 corpus.

3.2 RuSentiFrames Lexicon

The RuSentiFrames⁴ lexicon describes sentiments and connotations conveyed with a predicate in a verbal or nominal form (Rashkin et al., 2016; Klenner and Amsler, 2016). The structure of the frames includes the set of predicate-specific roles and frame dimensions.

For role designation, the approach of PropBank (Palmer et al., 2005) is used. In this approach, individual verb’s semantic arguments are numbered, beginning with zero. For a particular verb, Arg0 is generally the argument exhibiting features of a Prototypical Agent (Dowty, 1991), while Arg1 is a Prototypical Patient or Theme.

In the main part of the frame, the following dimensions are described:

- the attitude of the author of the text towards mentioned participants;
- positive or negative sentiment between participants;
- positive or negative effects to participants;
- positive or negative mental states of participants related to the described situation.

All assertions are provided with the score of confidence, which currently has two values: 1, if this assertion is true almost always, or 0.7, the assertion is considered in default. We do not describe assertions about neutral sentiment, effect or state of participants.

⁴<https://github.com/nicolay-r/RuSentiFrames/tree/v1.0>

Type of lexical unit	Number
Verbs	2 794
Nouns	822
Phrases	2 401
Other	49
Unique entries	6 036
Total entries	6 412

Table 2: Quantitative characteristics of the RuSentiFrames entries.

The created frames are associated not only with a single entry but with a "family" of related words and expressions, which have the same attitudes. The following lexical units can be associated with a sentiment frame: single words, idioms, light verb constructions, and some other multiword expressions.

Currently, RuSentiFrames contains 277 frames with 6,412 associated *frame entries*. Table 2 shows the distribution of the RuSentiFrames entries according to parts of speech (POS) and other characteristics. Let us consider frame "Одобрить" (Approve) presented in Example 1.

Example 1: Frame "Одобрить" (Approve)

```
"roles": {"a0": "who approves",
          "a1": "what is approved"}
"polarity": [{"a0", "a1", "pos", 1.0},
             ["a1", "a0", "pos", 0.7]},
"effect": [{"a1", "pos", 1.0}],
"state": [{"a0", "pos", 1.0},
          ["a1", "pos", 1.0]}
```

Nowadays, the lexicon is under development. For the proposed distant supervision approach, we utilize only the dimension of attitudes towards Prototypical Patient conveyed by Prototypical Agent. Table 3 provides related statistics.

Effect	Sentiment	Number
A0 → A1	pos	2 252
A0 → A1	neg	2 802

Table 3: The distribution of RuSentiFrames text entries according to attitudes.

3.3 News Collection

The collection to be used for sentiment attitude extraction consists of Russian articles and news of major news sources, specialized political sites, and

Russian sites of world known news agencies published in 2017.

Each article is separated into the title and the contents. The collection statistics presented in Table 4.

Parameter	Value
Number of documents	2.8×10^6
Sentences (avg./doc.)	13.24

Table 4: News collection statistics.

4 Automatic Forming of Training Collection for Sentiment Attitude Extraction

This section discusses two different methods of sentiment attitude annotation: pair-based, and frame-based. Both methods apply to the title as it provides the main idea of the article and usually has a relatively simple sentence structure. Figure 1 illustrates the collection development flow. Further subsections describe the flow components in detail.

4.1 Text Processing and Named Entity Recognition

For attitude extraction, it is necessary to parse a text. This process involves the *tokenization* to demarcate text string into words and punctuation signs. Numbers and URL-links are considered as non-meaningful and masked.

Each attitude is based on a pair of named entities. For named entity recognition (NER, Figure 1), we utilize the following resources:

1. The pre-trained neural network model⁵, which is state-of-the-art for the Russian language (Burtsev et al., 2018);
2. The list of named entities from RuSentRel corpus, organized in *authorized objects*.

The list of authorized objects is necessary to avoid accidental misses from the NER model.

The text of news articles may refer to an entity in several naming variants (Putin – Vladimir Putin), and synonyms (EU – Europe). To match named entity synonyms, in this paper we utilize both stemming⁶ and list of synonyms, provided along with the RuSentRel corpus.

⁵<https://github.com/deepmipt/ner>

⁶<https://tech.yandex.ru/mystem/>

4.2 Pair-Based Annotation

This attitude annotation method utilizes the pre-assigned attitudes organized in a *list of pairs* (Figure 1).

Given a processed title with labeled named entities set E , we select pairs $\{\langle e_i, e_j \rangle \mid e_i, e_j \in E\}$, suitable for sentiment attitudes role. For relevant pairs filtering, the following restrictions should be met:

1. The presence of synonymous attitude in a given attitudes list;
2. All the entities appeared between pair endings should be authorized objects;

However, in a specific sentence, the supposed relation between countries can be false. For example, in the sentence “Зрители смогут увидеть показательные выступления спортсменов - чемпионов России и Европы” (Spectators will be able to see demonstrations of athletes - champions of Russia and Europe), prior negative relations between Russia and Europe are not mentioned. Therefore we need an additional factor to provide the quality of the annotation, and the RuSentiFrames lexicon can be used as such a factor.

4.3 Frame-Based Annotation

This attitude annotation method utilizes frame entries from the RuSentiFrames lexicon. Given a processed title with labeled named entities set E , an entry pair with $e_i, e_j \in E$, where e_i appears before e_j , considered as sentiment attitude when all the following criteria are met:

- All the frame entries between e_i and e_j have polarity;
- All the entities that appeared between e_i and e_j should be authorized.

We assign a positive sentiment score when all the polarities of inner frame entries have a positive sentiment. Otherwise, it assigns the negative sentiment score. We also consider frame entry polarity as *inverted*, when it is used with “не” (not) particle.

4.4 Attitude Filtering

To combine the annotation methods described above (attitudes filter, Figure 1), we intersect the annotations and separate the intersection into the




Figure 1: Training collection development flow

following sets: (i) with the same polarity and (ii) with different polarity according to both sources.

In the case of the non-empty set with the same polarity (SAME), at last step we utilize *sentence filter* (Figure 1). Given a set of processed news sentences, we select those which contain at least a single entity pair, presented in the SAME set.

Corpus	doc. level attitudes	texts count	titles and sentences
Pair-Based	60 788	52 377	136 496
Frame-Based	55 566	43 383	104 205
Intersection	22 589	20 885	50 958
Different	7 929	7 435	17 939
Same _{RuAttitudes}	14 660	13 450	33 019

Table 5: The statistics of automated annotation of texts and sentences.

Finally, the workflow (Figure 1) is applied to the news collection (Section 3.3), and we obtain the **RuAttitudes**⁷ dataset, automatically labeled with sentiment attitudes between named entities. Table 5 provides statistics separately for each step. Evaluated accuracy of randomly selected sentences from different texts presented in Table 6.

5 Convolutional Neural Networks for Attitude Classification

For automatic sentiment attitude classification, the following CNN-based architectures were used:

⁷<https://github.com/nicolay-r/RuAttitudes/tree/v1.0>

Corpus	Accuracy
Pair-Based	67.0
Frame-Based	62.5
RuAttitudes	89.0

Table 6: Accuracy of attitude annotation in the generated collections.

- Classic CNN (Zeng et al., 2014);
- Piecewise-CNN (Zeng et al., 2015);

To predict the attitude polarity, both models utilize sentences in the input. Given a context (a set of sentences) with a mentioned attitude in it, the output of models is the class label. Two different approaches to training were considered: single-sentence training and multi-sentence training.

5.1 Sentence Embedding

We use *sentence embedding* to present sentences in model input. This is a matrix with rows related to words or token (for example, punctuation marks) embeddings.

For words, we look up for vectors in precomputed and publicly available Word2Vec model⁸ based on news articles with *window size* of $w = 20$. For tokens, we utilize a set of predefined types (size of 17), where each type is a randomly initialized vector of the same size as word vector.

Each word and token vectors have been additionally expanded with the following features:

⁸http://rusvectors.org/static/models/rusvectors2/news_mystem_skipgram_1000_20_2015.bin.gz

- Distance embedding (Rusnachenko and Loukachevitch, 2018) – is vectorized distance in words from entities e_i and e_j of an entry pair $\langle e_i, e_j \rangle$ to a given word or token;
- Part-of-speech (POS) tags embedding; we use “unknown” tag in case of tokens.

For features, we use randomly initialized vectors. Table 7 provides parameter values of each embedding described above.

Type	Parameters	Values
POS	v_{size}	5
Distance	v_{size}	5
Tokens	$\langle size, l_t \rangle$	$\langle 17, 10^3 \rangle$
Words	$\langle size, l_w, w \rangle$	$\langle 147 \cdot 10^3, 10^3, 20 \rangle$

Table 7: Embedding parameters, where v_{size} is the size of embedding vectors.

5.2 Single Sentence Training

This training process assumes to predict a sentiment label by a single sentence. Given an attitude context, we consider that each sentence should be labeled with an attitude sentiment.

We utilize training process described in (Rusnachenko and Loukachevitch, 2018). The input organized in minibatches, which yields n bags.

Each bag has a set of m sentences $\{s_1, \dots, s_m\}$, where $s_j = \langle e_s, y \rangle$ includes sentence embedding e_s and related label $y \in \mathbb{R}^c$. The training process is iterative, and iteration includes the following steps:

1. Composing a minibatch $I = \{b_1, \dots, b_n\}$ where $b_i = \{s_1, \dots, s_m\}$;
2. Performing a forward propagation through the network; the result is a vector $\{o_k\}_{k=1}^q$, where $o_k \in \mathbb{R}^c$, and $q = n \cdot m$;
3. Computing cross entropy loss for output:

$$l_k = \sum_{j=1}^c \log p(y_i | o_{k,j}; \theta), \quad k \in \overline{1 \dots q} \quad (1)$$

4. Composing cost vector $\{cost_i\}_{i=1}^n$, where $cost_i = \max [l_{i.g} \dots l_{(i+1).g}]$ is a maximal loss within i 'th bag;
5. Using $cost$ to update hidden variables set;

5.3 Multiple Sentence Training

This training process assumes to predict a sentiment label for sentences set. This process refers to a single sentence case described in Section 5.2 with the following modifications:

- Each minibatch presented as a sequence of n bags $b_i = \langle E_s, y \rangle$, where E_s is a set of embedded sentences (Section 5.2);
- Result output vector $\{o_l\}_{l=1}^n$ obtained by an application of max-pooling operation over separately convolved context sentences (Jiang et al., 2016);

6 Datasets and Experimental Setup

We consider the problem of sentiment attitudes classification as a two-class task at the document level. We conduct three different experiments using the following datasets:

1. RA – RuAttitudes dataset, described in this paper (Section 4);
2. RSR – RuSentRel based dataset with sentence-level attitude labeling (Section 3.1);
3. RSR+RA – a combination of RSR and RA datasets;

By default, the RuSentRel dataset provides a document level attitude labeling. This labeling was used to complete RSR and therefore treat RuSentRel in the same format as RuAttitudes. We consider sentence-level attitudes as bidirectional pairs of named entities. We filter sentences from RuSentRel with sentiment pairs according to the following rule. For each sentence among the all documents we check, whether at least a single pair has the labeled attitude, presented in document attitude annotation.

Statistical comparison of the RSR and RuAttitudes datasets is presented in Table 8. We use 10-fold cross-validation (CV) in RA experiment. In other experiments, the 3-fold CV has been chosen due to a small number of documents in RSR dataset. In experiments with RSR+RA, the cross-validation procedure applies to RSR; the RuAttitudes dataset combines with each training block of RSR. It is worth noting that RuSentRel sentences have a larger amount of opinions per sentences (2.26), mainly due to the nature of their content – these are analytical articles, while RuAttitudes has

been based on news reports. This fact makes experiments with the RSR dataset significantly challenging.

For models to be trained, we apply *named entity masking* for pairs due to: (i) omit entity-related feature dependency, and (ii) prevent models from learning the distribution of the latter.

Parameter	RuAttitudes	RSR
Documents	20 855	73
k -fold cross-validation	10	3
Opinions on sentence level	35 125	2 879
Negative opinion pairs	26 904	1 602
Positive opinion pairs	8 221	1 277
Avg. opinions per sentence	1.06	2.26
Avg. sentences per opinion	2.40	2.57

Table 8: Comparison of RuAttitudes and RuSen-tRel based (RSR) datasets for experiments.

Description	Parameters	Values
Minibatch	$\langle n, m \rangle$	$\langle 8, 3 \rangle$
Optimiser	$\langle lr, \rho, \epsilon \rangle$	$\langle 0.1, 0.95, 10^{-6} \rangle$
Terms	k	50
Window size	w	3
Filters count	c	300
Dropout	ρ	0.9

Table 9: Predefined training parameters.

Table 9 illustrates model parameter values. Each minibatch has n bags. As for the sentence count per bag parameter, we select $m = 3$ to cover the average sentence count per opinion (Table 8). To translate the labels onto the document level, we utilize *average* function across all the sentences (sentence sets) of a given attitude. All the sentences were limited by k words, including tokens. Each word is considered in a lemmatized⁹ form. The convolutional window size and the filters count were chosen according to (Zeng et al., 2015). We use the *adadelta* optimizer with parameters according to (Zeiler, 2012).

Several baselines were also added in experiments: **baseline_neg** – all pairs of named entities are labeled as negative; **baseline_rand** – pairs are labeled randomly according to the sentiment distribution in the training collection;

We measure average values of accuracy every five epochs. The training process terminates if one of the following conditions are met: (i) average

⁹<https://tech.yandex.ru/mystem/>

epoch accuracy reaches 99%; (ii) the training error exceeds the prior related value (except RSR experiment). The latter exception is due to an unstable training process, that might be caused by the relatively small training set passed per a single epoch.

For this task, we adopt macroaveraged (over documents): F1-score (F_1^{PN}), precision (π^{PN}), and recall (ρ^{PN}).

7 Result Analysis and Discussion

Table 10 provides the results of both evaluated baselines and models for each dataset. For RuAttitudes dataset, where the negative class significantly exceeds positive, and most documents lack positive attitudes, both baselines show high values.

For a certain model and related experiments α and β , let $C_{\alpha,\beta}$ is a set of labeling contradictions of two experiments. All the pairs of contradictions could be then treated as correctly or wrongly labeled by each experiment. We let $C_{\alpha,\beta}^*$ as a subset of contradictions where β corrects the errors of α . This subset yields of only those pairs, where the sentiment class has been correctly defined in β . Table 11 provides the comparison statistics between RSR and RSR+RA experiments, separately for each model, where $|L|$ is an opinion pairs count (Table 1). While the contradiction represents 34% of total opinion pairs, with an average 61% of corrections and 38% of wrong labeling, we can conclude the average result error corrections are 7% in case of RSR+RA experiment (Table 10).

The contribution of RuAttitudes corpus in RSR+RA experiment could be considered in terms of frequencies of following entry types E : nouns, verbs, frames. Due to the task considered as sentiment attitudes classification, it is significant to separate statistics by positive and negative classes. To define a sentence class, we utilize sentiment of attitude that appears in it. We utilize semantic orientation (SO) function (Turney, 2002) to reveal a discrepancy in entries between sentiment classes:

$$SO(e) = PMI(e, pos) - PMI(e, neg) \quad (2)$$

where $PMI(e, c)$ is a *pointwise mutual information* of entry e and sentiment class c . For each entry type separately, we utilize Formula 2 towards RuAttitudes to complete a set of entries, bounded with positive ($SO(e) > 0$) and negative ($SO(e) < 0$) classes. We order these sets by descending of $|SO(e)|$ to select k most distinctive entries and complete RA_k^c subsets, $c \in \{pos, neg\}$.

Models	RA			RSR			RSR+RA		
	F_1^{PN}	π^{PN}	ρ^{PN}	F_1^{PN}	π^{PN}	ρ^{PN}	F_1^{PN}	π^{PN}	ρ^{PN}
baseline-neg	0.83	0.78	0.89	0.39	0.31	0.54	0.39	0.31	0.54
baseline-rand	0.72	0.75	0.71	0.49	0.51	0.48	0.49	0.51	0.48
CNN	0.91	0.88	0.94	0.52	0.52	0.55	0.63	0.62	0.66
PCNN	0.93	0.91	0.96	0.59	0.58	0.61	0.67	0.66	0.69
MI-CNN	0.91	0.88	0.94	0.57	0.56	0.60	0.62	0.60	0.65
MI-PCNN	0.94	0.92	0.96	0.62	0.60	0.64	0.68	0.67	0.70

Table 10: Result of single sentence (CNN, PCNN) and multiple sentence (MI-CNN, MI-PCNN) trained models in following experiments: RA – RA results trained on RA; RSR – RSR results trained on RSR; RSR+RA – RSR results trained on RSR+RA.

Model	$ C_{2,3} $	$ L / C_{2,3} $	$ C_{2,3}^* / C_{2,3} $
CNN	468	0.34	0.63
PCNN	428	0.31	0.62
MI-CNN	488	0.36	0.58
MI-PCNN	442	0.35	0.60

Table 11: Contradiction statistics between experiments RSR (2) and RSR+RA (3).

To assess how RuAttitudes effects on error corrections, we provide statistic of entries both appears in $C_{2,3}^*$ and RA_k^c . For each entry $e \in RA_k^c$ we calculate $tf(e, C_{2,3}^*)$ – is an averaged (among all models) normalized term frequency of entry e in $C_{2,3}^*$. Table 12 lists three ($k = 3$) most distinctive entries by each entry type, where entries with $tf(e, C_{2,3}^*) > 0.5$ are bolded. It is possible to investigate that $C_{2,3}^*$ mostly saturated with positively bounded frames of RA_k^{pos} and negatively bounded nouns of RA_k^{neg} .

Conclusion

This paper proposes an approach to the automatic development of a train collection for the sentiment attitude extraction task in the news domain. The combination of two different techniques was used to provide the double-check and keep labeled results of the common intersection. The first proposed technique obtains contexts by a manually implemented list of pairs of named entities with sentiment scores. The other technique, on the contrary, extracts relations from contexts using sentiment frames. The latter became possible due to the assumption of title simplicity.

Sentiment attitude extraction was considered as a two-class classification task. This result analysis demonstrates the model classification improvements achieve 13.4% increase in F_1^{PN} when the

E	Entry Value	$tf(e)$
N_{pos}	«поддержка» (support)	0.20
	«помощь» (help)	0.02
	«переговоры» (negotiations)	0.24
V_{pos}	«поддерживать» (to support)	0.55
	«начинать» (to start)	0.26
	«предлагать» (to suggest)	0.10
F_{pos}	«помочь» (to help)	0.23
	«начать» (to begin)	0.68
	«договориться» (to agree)	0.98
N_{neg}	«санкция» (sanction)	0.78
	«борьба» (fight)	0.50
	«отношение» (relation)	0.96
V_{neg}	«обвинять» (to blame)	0.05
	«вводить» (to introduce)	0.06
	«продлять» (to extend)	0.00
F_{neg}	«наказать» (to punish)	0.59
	«обвинить» (to blame)	0.06
	«бороться» (to fight)	0.44

Table 12: List of $k = 3$ most distinctive nouns (N), verbs (V), and frames (F) of RA_k^c with related frequencies in $C_{2,3}^*$; E is an entry type.

latter being trained with the developed collection.

In further work, we plan to address the shortcomings in the following directions: to emphasize the difference between sentiment and non-sentiment relations and to reduce noisy labeling of the existed approach.

Acknowledgments

The reported study was funded by RFBR according to the research project № 19-37-50001. The development of Russian sentiment frames is supported by the RFBR research project № 16-29-09606.

References

- Mikhail Burtsev, Alexander Seliverstov, Rafael Airapetyan, Mikhail Arkhipov, Dilyara Baymurzina, Nickolay Bushkov, Olga Gureenkova, Taras Khakhulin, Yuri Kuratov, Denis Kuznetsov, et al. 2018. Deeppavlov: Open-source library for dialogue systems. In *Proceedings of ACL 2018, System Demonstrations*. pages 122–127.
- Eunsol Choi, Hannah Rashkin, Luke Zettlemoyer, and Yejin Choi. 2016. Document-level sentiment inference with social, faction, and discourse context. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. volume 1, pages 333–343.
- Lingjia Deng and Janyce Wiebe. 2015a. Mpqa 3.0: An entity/event-level sentiment corpus. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. pages 1323–1328.
- Lingjia Deng and Janyce Wiebe. 2015b. Mpqa 3.0: An entity/event-level sentiment corpus. *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* pages 1323–1328.
- David Dowty. 1991. Thematic proto-roles and argument selection. *language* 67(3):547–619.
- Joe Ellis, Jeremy Getman, and Stephanie M Strassel. 2014. Overview of linguistic resources for the tac kbp 2014 evaluations: Planning, execution, and results. In *Proceedings of TAC KBP 2014 Workshop, National Institute of Standards and Technology*. pages 17–18.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2009. Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. In *Proceedings of the Workshop on Semantic Evaluations: Recent Achievements and Future Directions*. Association for Computational Linguistics, pages 94–99.
- Xiaotian Jiang, Quan Wang, Peng Li, and Bin Wang. 2016. Relation extraction with multi-instance multi-label convolutional neural networks. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*. pages 1471–1480.
- Manfred Klenner and Michael Amsler. 2016. Sentiframes: A resource for verb-centered german sentiment inference. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*. pages 2888–2891.
- Natalia Loukachevitch and Nicolay Rusnachenko. 2018. Extracting sentiment attitudes from analytical texts. *Proceedings of International Conference on Computational Linguistics and Intellectual Technologies Dialogue-2018 (arXiv:1808.08932)* pages 459–468.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*. Association for Computational Linguistics, pages 1003–1011.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational linguistics* 31(1):71–106.
- Hannah Rashkin, Sameer Singh, and Yejin Choi. 2016. Connotation frames: A data-driven investigation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. pages 311–321.
- Nicolay Rusnachenko and Natalia Loukachevitch. 2018. Neural network approach for extracting aggregated opinions from analytical articles. In *International Conference on Data Analytics and Management in Data Intensive Domains*. Springer, pages 167–179.
- Peter D Turney. 2002. Thumbs up or thumbs down?: semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th annual meeting on association for computational linguistics*. Association for Computational Linguistics, pages 417–424.
- Matthew D Zeiler. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701* .
- Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. pages 1753–1762.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*. pages 2335–2344.