

June 14–16, 2023

RuSentNE-2023: Evaluating Entity-Oriented Sentiment Analysis on Russian News Texts

Anton Golubev
Lomonosov Moscow
State University
antongolubev5@yandex.ru

Nicolay Rusnachenko
Bauman Moscow State
Technical University
rusnicolay@gmail.com

Natalia Loukachevitch
Lomonosov Moscow
State University
louk_nat@mail.ru

Abstract

The paper describes the RuSentNE-2023 evaluation devoted to targeted sentiment analysis in Russian news texts. The task is to predict sentiment towards a named entity in a single sentence. The dataset for RuSentNE-2023 evaluation is based on the Russian news corpus RuSentNE having rich sentiment-related annotation. The corpus is annotated with named entities and sentiments towards these entities, along with related effects and emotional states. The evaluation was organized using the CodaLab competition framework. The main evaluation measure was macro-averaged measure of positive and negative classes. The best results achieved were of 66% Macro F-measure (Positive+Negative classes). We also tested ChatGPT on the test set from our evaluation and found that the zero-shot answers provided by ChatGPT reached 60% of the F-measure, which corresponds to 4th place in the evaluation. ChatGPT also provided detailed explanations of its conclusion. This can be considered as quite high for zero-shot application.

Keywords: Targeted Sentiment Analysis, Named Entity, News Texts, ChatGPT

DOI: 10.28995/2075-7182-2023-22-130-141

RuSentNE-2023: анализ тональности по отношению к сущностям в русскоязычных новостных текстах

Голубев А. А.
МГУ им. Ломоносова
Москва, Россия
antongolubev5@yandex.ru

Русначенко Н. Л.
МГТУ им. Баумана
Москва, Россия
rusnicolay@gmail.com

Лукашевич Н. В.
МГУ им. Ломоносова
Москва, Россия
louk_nat@mail.ru

Аннотация

В статье описывается тестирование RuSentNE-2023, посвященное таргетированному анализу тональности в русскоязычных новостных текстах. Задача участников состояла в том, чтобы предсказать тональность по отношению к именованному объекту в предложении. Датасет тестирования RuSentNE-2023 основан на корпусе российских новостей RuSentNE, в котором размечено несколько типов явлений, связанных с тональностью. Корпус аннотирован именованными сущностями, тональностью по отношению к этим сущностям, размечены последствия для сущностей в связи описываемыми ситуациями и эмоциональные состояниями сущностей. Тестирование было организовано на основе специализированного сайта для тестирований CodaLab. Основной мерой оценки было макроусреднение положительных и отрицательных классов. Наилучший результат, полученный участниками, был 66% макро-F-мере (положительные + отрицательные классы). Мы также протестировали ChatGPT на тестовом наборе нашего тестирования и обнаружили, что zero-shot (без обучения) ответы ChatGPT достигают 60% F-меры, что соответствует 4-му месту в тестировании RuSentNE. Модель ChatGPT также предоставила подробные пояснения к своему заключению. Этот результат можно считать достаточно высоким для применения в условиях zero-shot.

Ключевые слова: таргетированный анализ тональности, именованная сущность, новостные тексты, ChatGPT

1 Introduction

Sentiment analysis studies began with the general task setting, in which the general sentiment of a sentence or a text should be detected. Currently, so-called targeted sentiment analysis is intensively discussed, in which a model should determine the attitude towards specific entities, their aspects (properties), or topics.

Targeted sentiment analysis is especially important for news flow processing. It is assumed that news texts should be neutral, but in fact they contain a variety of information about positions on various issues of state bodies, companies, opinions of individuals, positive or negative attitudes of the mentioned subjects to each other. All these opinions are important for collecting and analysis.

Currently, news sentiment seems understudied. For example, the search on paper titles in Google Scholar shows that number of papers devoted to sentiment in social media is three times larger than the number of paper discussing news sentiment.

The specific features of news texts from the point of view of sentiment analysis are as follows (Loukachevitch and Rusnachenko, 2018):

- these texts contain various opinions conveyed by different subjects, including the author(s)' attitudes, positions of cited sources, and relations of mentioned entities to each other;
- some sentences are full of named entities with different sentiments, which makes it difficult to determine sentiment for a specific named entity;
- news texts contain numerous named entities with neutral sentiment, which means that the neutral class largely dominates;
- significant share of sentiment in news texts is implicit, for example can be inferred from some actions of entities.

The authors of (Hamborg et al., 2021) annotated a news corpus with sentiment and stated that sentiment in the news is less explicit, more dependent on context and the reader's position, and it requires a greater degree of interpretation. It was found that state-of-the-art approaches to targeted sentiment analysis perform worse on news articles than in other domains. The authors also point out that in 3% of cases, the sentiment is depends on the position of the reader.

In this paper we present the Russian News corpus RuSentNE annotated with named entities and sentiment towards these entities, related effects and emotional states. We used the sentiment annotation of the RuSentNE corpus to organize the shared task RuSentNE-2023 within the framework of Dialogue evaluation series.

2 Related Work

There were several international evaluations devoted to targeted sentiment analysis.

In 2012-2014 within the CLEF conference, RepLab events devoted to the evaluation of online reputation management systems were organized (Amigó et al., 2012; Amigó et al., 2014). The task was to determine if the tweet content has positive or negative implications for the company's reputation.

In the SemEval evaluation workshops 2015, 2016, studies were devoted to aspect-based sentiment analysis (ABSA) in several domains. The task was to determine sentiment towards specific characteristics discussed in users' reviews such as food or service in restaurants (Pontiki et al., 2015; Pontiki et al., 2016). In 2016-2017, topic-oriented sentiment analysis tasks were studied in the SemEval series of Twitter sentiment evaluations (Nakov et al., 2016; Rosenthal et al., 2017).

The latest trend in targeted sentiment analysis is the so-called Structured Sentiment Analysis, which involves extracting tuples from texts that describe opinions of the following form $\langle h, t, e, p \rangle$, where h is the opinion holder, p represents sentiment (positive or negative), in relation to the entity t , expressed by means of the word or phrase e . Structured sentiment analysis can be divided into five subtasks: i) sentiment expression extraction, ii) sentiment object extraction, iii) sentiment subject extraction, iv) determination of the relationship between these elements, and v) sentiment extraction (positive or negative). Modern approaches aim to address these problems in a unified manner, generating the required tuples (Lin et al., 2022). In 2022, the competition on structural sentiment analysis competition was organized as part of the SemEval-2022 international evaluation workshop (Barnes et al., 2022).

The relatively recent advent of *transformers* (Vaswani et al., 2017) cause a significant breakthrough in machine learning application across the variety of natural language processing tasks, including targeted sentiment analysis. Within the last five years, the significant amount of works were aimed on application of transformer components, namely the encoder (Devlin et al., 2019) and decoder (Alt et al., 2019a). In the case of the application of BERT, complemented by classification layer on top, has resulted in the appearance of a variety pretrained models employing different pretrain techniques (Zhuang et al., 2021; Alt et al., 2019b). Another group of authors studies conveying the structure of target in context, emphasizing the target in texts and forming a prompt-based input constructions (Sun et al., 2019a; Shin et al., 2020). To the best of our knowledge, the structured representation of the input contexts (Morio et al., 2022) as a part of sequence-to-sequence and graph-based models represents the latest advances in target oriented sentiment analysis.

For Russian, targeted sentiment evaluations were organized as a two-year series in 2015 and 2016. In 2015, aspect-based sentiment evaluation in restaurant and car reviews was explored (Loukachevitch et al., 2015). During two years, methods for reputation monitoring tasks towards banks and telecommunication companies in tweets were evaluated (Loukachevitch and Rubtsova, 2015). The best classification results in 2015-2016 competitions were mainly based on the SVM method and the GRU neural network. Later, the results were significantly improved with the application of the BERT model (Golubev and Loukachevitch, 2020; Golubev and Loukachevitch, 2021).

We see that most international and Russian evaluations are devoted to targeted sentiment analysis applied to social media and user reviews, not to news analysis. The general task of sentiment analysis devoted to classification of Russian news quotes was studied in the ROMIP-2013 evaluation (Chetviorkin and Loukachevitch, 2013).

The work closest to ours is the NewsTSC corpus (Hamborg et al., 2021). The authors selected articles of 14 US newspapers from the Common Crawl news crawl (CC-NEWS). Mentions of named entities such as PERSON or ORG were automatically identified, and corresponding sentences were extracted. The annotators should read the sentence and determine sentiment towards a highlighted named entity. In total, 3002 sentences were annotated, the neutral sentiment was most frequent (2087). Several BERT models were used in the experiments. The best results (58.8 macro F measure) were achieved LCF-BERT (Zeng et al., 2019), where BERT was tuned on the Common Crawl news corpus.

3 RuSentNE Sentiment-Annotated Corpus

The dataset RuSentNE-2023 constructed for the shared task, is based on the RuSentNE corpus, in which diverse sentiment-related phenomena are annotated.

The corpus RuSentNE includes news texts from the NEREL dataset, annotated with 29 named entity types (Loukachevitch et al., 2021). 400 NEREL texts with the largest relative share of sentiment words according to the RuSentiLex lexicon (Loukachevitch and Levchik, 2016) were selected for sentiment annotation in RuSentNE. The annotated named entities were used as targets of sentiment in the RuSentNE corpus.

Then sentiment-related tags and relations towards NEREL entities were labeled in the RuSentNE. The sentiment annotation includes 12 entity tags and 11 relation types, which describe sentiment, arguments, effects and emotions of entities mentioned in the text.




Figure 1: Example of sentiment annotation in sentence "Apple and Samsung infringed each other's patents" shows negative relations between companies and negative argument towards them

Figures 1, 2 and 3 show examples of annotation in the RuSentNE corpus. In Figure 1, we can see that two companies are negative to each other (NEGATIVE_TO relation); also negative argument ("enfringed patents") is annotated that explains the attitude of two companies towards each other

(OPINION_RELATES_TO relation).

In Figure 2, we see that the Matteo Renzi is positive to Fo (POSTITIVE_TO relation); the explanation for this attitude is given (*major figure in cultural life*). The emotional state of Matteo Renzi is negative because of the Fo death. Also negative effect for Dario Fo is annotated, which originated from his death. In Figure 3 we see the author position towards Berlusconi depicted as tag AUTHOR_NEG.




Figure 2: Example of sentiment annotation in sentence "Prime Minister of Italy Matteo Renzi said that with Fo's departure, "the country has lost a major figure in cultural life." demonstrates positive attitude from Matteo Renzi towards Dario Fo and also negative emotions stemming from his death.




Figure 3: Example of sentiment annotation in sentence "Notorious figure in the country - Berlusconi has been repeatedly accused of financial fraud." shows negative opinion of the author.

For the RuSentNE annotation, the BRAT annotation tool was used. The texts were annotated by both students and specialists in computational linguistics. All texts were then checked by a moderator, who discussed any identified issues' with annotators.

4 RuSentNE-2023 Dataset

The annotation of the RuSentNE corpus was partially used for the RuSentNE-2023 evaluation. Since RuSentNE contains 29 different classes of named entities, we selected those entities that are most frequent objects of the targeted sentiment in news texts. We selected the following classes of entities for the evaluation:

- PERSON — physical person regarded as an individual,
- ORGANIZATION — an organized group of people or company,
- COUNTRY — a nation or a body of land with one government,
- PROFESSION — jobs, positions in various organizations, and professional titles,
- NATIONALITY — nouns denoting country citizens and adjectives corresponding to nations in contexts different from authority-related.

The distribution of entity types in the training, validation, and test sets is presented in Table 1

Stage	Total	PERSON	ORG.	COUNTRY	PROF.	NATIONAL.
Train	6637	1934 (29%)	1487 (23%)	1274 (19%)	1666 (25%)	276 (4%)
Development	2845	857 (30%)	653 (23%)	686 (19%)	533 (24%)	116 (4%)
Final	1947	480 (25%)	484 (25%)	363 (19%)	510 (26%)	110 (5%)

Table 1: Distribution of entity types in training, validation and test sets

For RuSentNE-2023 evaluation devoted to targeted sentiment analysis, we used a subset of annotated entity tags and relations from the initial annotation of the RuSentNE corpus:

- AUTHOR_POS, AUTHOR_NEG tags describe the author's attitude towards the tagged entity,
- POSITIVE_TO, NEGATIVE_TO relations describe a relationship between two entities in a text,
- OPINION_RELATES_TO relation describes the relation from the opinion expressed in the text to the object of the opinion.
 - OPINION_WORD_NEG, ARGUMENT_NEG — negative to the object,

– OPINION_WORD_POS, ARGUMENT_POS — positive to the object

Entities without any sentiment annotations or relations directed towards them form the neutral class.

In particular, the following examples of targeted sentiment in the NEREL-2023 dataset are obtained from the annotation presented in Figures 1, 2, and 3 (Table 2).

Source	Entity	Sentiment	Source in RuSentNE
Figure 1	Samsung	negative	NEGATIVE_TO, ARGUMENT_NEG
Figure 1	Apple	negative	NEGATIVE_TO, ARGUMENT_NEG
Figure 2	Fo	positive	POSITIVE_TO, ARGUMENT_POS
Figure 2	Matteo Renzi	neutral	absense of annotation
Figure 3	Berlusconi	negative	AUTHOR_NEG, ARGUMENT_NEG

Table 2: Examples of sentiment labels in RuSentNE-2023 dataset

In the original RuSentNE corpus, sentiment-related relations can be annotated across sentences. In the RuSentNE-2023 evaluation, the targeted sentiment should be extracted from a single sentence. This implies that in a specific sentence, sentiment towards a target entity can be absent. Therefore, it was then necessary to ensure that the source of sentiment is located in the same sentence as the entity. If not, the sample was not included in the dataset. In the example below, the Tula transport prosecutor’s office from the second sentence derives the positive annotation from the first one, which contradicts the relations from isolated second sentence:

- «*The Tula transport prosecutor’s office **defended the rights** of workers. The Tula transport prosecutor’s office **filed two lawsuits against Russian Railways** for the recovery of child care allowances*».

The pre-trained RuCoreNews_{small} spaCy model¹ was utilized to segment texts into sentences.

The initial RuSentNE corpus contains so-called nested named entities, that is one entity can be annotated within another entity. Only an upper-level entity was selected for the RuSentNE-2023 dataset. An example of nested entities can be seen in Figure 2: entity *Prime-minister of Italy* contains entity *Italy*.

After the collection was formed, several post-processing steps were carried out. The same sentences that were included in the dataset multiple times according to different criteria were deduplicated. Examples with conflicting annotations were excluded. The minimum (40) and maximum (430) lengths of the text were chosen experimentally. Sentences excluded at this stage are presented below. The first sentence contains conflicting annotation for Iran (Tel Aviv vs Syria, Hezbollah and Hamas), while the following two sentences are too short:

- «*Tel Aviv believes that the price of the withdrawal of Israeli troops from the Golan Heights should be political concessions from Syria - the weakening of its ties with its strategic ally **Iran** and the cessation of support for the Lebanese resistance movement Hezbollah and the Palestinian Hamas*».
- «*The **shepherd** was not injured*».
- «***Pirates** fought back*».

Table 3 presents the distribution of sentiment scores in the training, validation and test sets.

Type	Stage	Total	Positive	Negative	Neutral
train	Train	6637	856 (13%)	1007 (15%)	4774 (72%)
validation	Development	2845	362 (13%)	438 (15%)	2045 (72%)
test	Final	1947	269 (14%)	243 (12%)	1435 (74%)

Table 3: Distribution of sentiment scores in training, validation and test sets.

We estimated the inter-annotated agreement in the RuSentNE-2023 dataset using duplicating sentences from the initial dataset. Cohen’s Kappa is calculated as 0.5 (moderate agreement). But the Cohen’s Kappa can be unreliable in our case due to Kappa’s sensitivity to class imbalance (Feinstein and Cicchetti, 1990). The percentage agreement between annotators is 84%.

¹https://spacy.io/models/ru#ru_core_news_sm

5 Task Description

In the RuSentNE-2023 competition, named entities should be classified into one of three sentiment classes: positive, negative or neutral within the context of a single sentence. Each sentence is annotated as follows:

- `entity` — object of sentiment analysis
- `entity_tag` — tag of this object (see Section 4)
- `entity_pos_start_rel` — index of the initial character of the given entity
- `entity_pos_end_rel` — index of the next character after the last of the given entity
- `label` — sentiment label

Each entity has a three-scaled label. The following classes (labels) are used:

- Negative (-1)
- Neutral (0)
- Positive (1)

Participants are tasked with automatically annotating each test sentence with the appropriate sentiment label for a specified entity.

As the primary evaluation metric, we adopt the F1-PN-macro, which averages the F1-measures of the positive and negative classes. Additionally, we calculated the traditional F1-macro measure as a supplementary metric.

6 Results

Over 15 participants took part in development stage of the competition. Table 4 illustrates the results obtained by the competitors during this preliminary evaluation stage.

Participant	F1-PN-macro (rank)	F1-PN0-macro (rank)
mtsai	70.94 (1)	77.63 (1)
cookies	69.89 (2)	76.74 (2)
lsanochkin	68.11 (3)	75.36 (3)
Dmitry315	62.91 (4)	71.28 (4)
ryzhtus	62.34 (5)	70.57 (6)
mitrokosta	62.15 (6)	70.70 (5)
sag_m	61.35 (7)	69.73 (7)
shershulya	60.07 (8)	69.26 (8)
s231644	59.99 (9)	69.10 (9)
antongolubev	57.73 (10)	68.21 (10)
ild	57.46 (11)	67.52 (11)
GreatDispersion	57.00 (12)	65.38 (12)
abc111	55.25 (13)	65.00 (13)
baseline_model	44.32 (14)	57.89 (14)
postoevie	41.09 (15)	48.28 (16)
angyling	35.37 (16)	43.76 (17)
AlexSMSU	31.94 (17)	49.25 (15)

Table 4: Results of the Development evaluation stage, ordered by F1-PN-macro; participant nicknames with top-3 results are bolded; `baseline_model` results correspond to the application of the baseline model

The results of the final evaluation stage are illustrated in Table 5. Participants were not limited in the number of submissions due to the relatively high threshold provided for both evaluation stages with 1K submissions limit. During the final stage, those participants who did not exceed the baseline result chose not to include their submissions in the leaderboard. In the following, we provide brief descriptions of the methods adopted by participants during the final evaluation stage.

baseline_model. DeepPavlov (Burtsev et al., 2018) pre-trained conversational RuBERT_{base}² is used. The model is fine-tuned treating the targeted sentiment analysis task as a question-answering problem (Sun et al., 2019b).

Participant	F1-PN-macro (rank)	F1-PN0-macro (rank)
mtsai	66.67 (1)	74.11 (2)
cookies	66.64 (2)	74.29 (1)
lsanochkin	62.92 (3)	71.20 (3)
ChatGPT **	60.06 (-)	70.79 (-)
antongolubev	59.64 (4)	69.04 (4)
sag_m	59.33 (5)	68.71 (5)
mitrokosta	58.68 (6)	67.54 (6)
Dmitry315	53.60 (7)	62.92 (7)
ild *	53.20 (-)	63.78 (-)
abc111	49.98 (8)	61.32 (8)
Naumov_al	46.96 (9)	54.92 (10)
baseline_model	40.92 (10)	56.71 (9)

Table 5: Results of the Final evaluation stage, ordered by F1-PN-macro; participant nicknames with top-3 results are bolded; «*» corresponds to post-evaluation submissions, made rightafter the end of the final stage; baseline_model results correspond to the application of the baseline model; «**» represents zero-shot answers of ChatGPT application baseline

mtsai. The participant experimented with the following the following models: RuRoBERTa_{large}³, XLM-RoBERTa_{large}⁴, RemBERT⁵ (Chung et al., 2021). Two separate transformers-based models were adopted: one model was used for the original input, while the other was applied to the input with masked entities. The participant utilized the predefined «[MASK]» token for entities, identified with the Named Entity Recognition (NER) approach. Weights ranging from 0 to 1.0 were applied to the output, with the sentiment classes assigned a weight of 1.0 and the neutral class assigned a weight of 0.1. The participant also implemented threshold for neutral class: when the neutral class has the highest probability, but its value below threshold, they select the most probable class among “positive” and “negative”. Participant employed an ensembling technique, which was based on a five-fold split of the dataset from the development stage. This technique involved training different transformers on various splits of data. In total, five models were used to ensemble the output.

cookies. The participant experiment with the following set of language models: language models RuBERT_{base}, RuBERT_{large}, RuRoBERTa_{large}. In terms of the input representation, various masks and markers for entity representation were considered. During the development phase, participant concluded that RuRoBERTa_{large} illustrates the highest results across the other models of experiment set, as well as the absence of difference across different entity representation formats. For the final version of the input representation format, participant decided to mark the entity within the input sentence. The best submission for the development stage represents a system based on RuRoBERTa_{large}, fine-tuned using the inverse probability weighting technique. In the final evaluation stage, participant also employed ensemble learning and implemented automatic annotation of the development set.

lsanochkin. The participant used a language model with a prompting technique according to the following assumption: passing entity (e) jointly with the sentence⁶ (T) in the following format: «[CLS] T [SEP] e ». The model was fine-tuned on the input text format with transformer-based classification pipeline (fully-connected layer on the top of a transformer model, cross-entropy loss). In terms

²<https://docs.deeppavlov.ai/en/master/features/models/bert.html>

³<https://huggingface.co/sberbank-ai/RuRoberta-large>

⁴<https://huggingface.co/xlm-roberta-large>

⁵<https://huggingface.co/google/rembert>

⁶like question answering or natural language inference prompt (Sun et al., 2019a)

of the embedding pooling for the sentiment class identification, the participant consider « [CLS] » token as a sentence representation.

antongolubev. The participant utilized the DeepPavlov RuBERT-base model fine-tuned in natural language inference (NLI) problem setting together with transfer learning approach. The model was pre-trained on targeted sentiment analysis data from SentiRuEval 2015-2016 evaluations (Loukachevitch et al., 2015) and automatically generated data from Russian news corpus (Golubev and Loukachevitch, 2021). In addition, the SMOTE (Chawla et al., 2002) augmentation approach for increasing of positive and negative classes was held.

sag_m. The participant's approach was based on the text-to-text generation approach using the ruT5_{large}⁷ model. The generated output text is one of the possible sentiment labels for the analyzed named entity: "negative", "neutral" or "positive". The application process involved experimenting with several variants of data preparation for the model input and applying output token filtering to derive the final class label. The best results were achieved using additional data for training: the CABSAR dataset⁸.

mitrokosta. The participant applied the RuBERT model. Entity token hidden states of the last layer were pooled. In addition, the hidden states of all layers were processed by convolutions and those that relate to the entity were averaged again. The latter yielded some improvements in quality.

ild. The participant adopts SBERT_{large} language model, pre-trained initially on financial news dataset⁹.

Dmitry315. The participant used the RuBERT model for fine-tuning on the provided data, with pooling for embeddings related to the entity.

Naumov_al. The participant considers the task as a multiclass-classification problem and experimented with Interactive Attention Network (IAN) (Ma et al., 2017). Numerous experiments were conducted using different word vector representations¹⁰. The participant experimented with several parameters including hidden_dim in the LSTM layers, the learning rate, and batch size. Alongside the competition dataset, the participant used CABSAR corpus¹¹, which contains Russian-language sentences from various sources: posts of the Live Journal social network, texts of the online news agency¹², and Twitter microblog posts.

We observe that in all approaches exceeding the baseline model results, participants used neural language models. Analyzing the results, it is worth to conclude the importance of the following findings, related to language model application (Vaswani et al., 2017):

1. a pretraining technique plays role: the contribution of RoBERTa results in higher performance rather than original BERT for the same sizes language models.
2. scale of the models is another direction: larger size of the model likely results in a higher baseline in the case of the fixed batch size and training evaluation methodology across all the model's sizes under consideration.
3. application of ensembling techniques allows gaining higher prediction results (the two first approaches in the final leaderboard).
4. task abstraction: this is a particular case when entities might be masked, marked or prompted which may assist with reaching the desired outcomes.

7 ChatGPT in RuSentNE-2023 evaluation

In (Zhang et al., 2022) authors report state-of-the-art results in the stance detection domain with zero-shot application of ChatGPT. The latter became a source of our inspiration to contribute with the related findings by analyzing responses for RuSentNE-2023 dataset.

We applied a conversational system that uses the GPT-3.5 model (Zhang et al., 2022)¹³, which comes

⁷<https://huggingface.co/ai-forever/ruT5-large>

⁸<https://github.com/sag111/CABSAR>

⁹https://huggingface.co/chrommium/sbert_large-finetuned-sent_in_news_sents

¹⁰ELMo, RuBERT_{large} and XLM-RoBERTa_{large} (the latter was used in final model), combined with the original IAN model

¹¹<https://github.com/sag111/CABSAR>

¹²lenta.ru

¹³<https://platform.openai.com/docs/models/gpt-3-5>

with a pretrained state¹⁴ corresponding to InstructGPT model. Access to this model is provided by the paid subscription service, ChatGPTPlus. The hidden state of the model was trained on text collections up to June 2021, and it supports up to 4,097 tokens inputs. The model is non-deterministic, implying that identical inputs can result in different outputs. We used a default temperature parameter. The sentences from the dataset, translated using the `googletrans` library¹⁵, were input into the ChatGPT model and were also incorporated into a prompt formatted as follows:

«What is the attitude of the sentence [translated sentence] to the target [translated entity]? Select one from «favor, against or neutral» and explain why».

Regular expressions with later manual validation were used to analyse model responses. For several examples (< 0.1% of data), the entity was translated incorrectly, which led to the fact that the model denied entity to be in the sentence. In such cases, the neutral label was put. Due to the limit on the number of requests to ChatGPT model per hour, a shell script was written to send examples at a given frequency. The whole test dataset was processed in 55 hours.

The results of ChatGPT are included in Table 5. We can see that the model applied in zero-shot manner (without fine-tuning) took the fourth place in the evaluation reaching more than 60% F1-PN-macro.

8 Analysis of Examples

To analyze difficult cases, we extracted examples from the test set, which were erroneously classified by all (9) or almost all (8) participants. We encountered the following main problem cases (the examples below are translated from Russian):

1. Models fail to distinguish between the subject and the object of opinion. For example, in the sentence "In 2011, Azerbaijan will increase pressure on international organizations in connection with the Nagorno-Karabakh problem", eight models inaccurately predicted negative sentiment towards Azerbaijan. However, according to the ChatGPT explanation, which seems correct in this context: "The sentence simply states a fact about the intentions of the Azerbaijani government without expressing any positive or negative opinion about it".

A similar issue can be seen in the sentence: "In it, Crowley announced that the Obama administration intends to continue to cooperate with the Israeli politician". Here most models inaccurately predicted positive sentiment towards the Obama administration, while both the annotators and ChatGPT indicated neutral sentiment.

2. Models do not distinguish some sentiment adjectives (evident for a human reader, but ambiguous) and fail to identify sentiment directed towards the target entity. For example, in sentence "In the very first days of Moiseev's stay at the clinic of JSC "Medicina", the patient was examined by the country's **leading** neurosurgeon, who did not find any indication for surgery", nearly all models, including ChatGPT, failed to detect positive sentiment towards "neurosurgeon".

3. Models fail to identify sentiment that requires an understanding of the relationships between entities. For example, in sentence "On the first day of 2011, numerous messages appeared on the **Microsoft** forum from Hotmail users complaining about the disappearance of all read messages stored in their mailboxes". In this case, all models (including ChatGPT) did not detect the negative sentiment towards Microsoft, the owner of Hotmail service.

4. Models do not distinguish implicit sentiment. For example, in sentence "Over the years of her film career, she has appeared **in more than 30 films and television series**, including the sci-fi film *Forbidden Planet* and the detective series "Honey West", for which the actress **received a Golden Globe film award** and was **nominated for an Emmy Award.**" both the models and ChatGPT were unable to detect the positive sentiment toward the actress.

We also encountered several issues with human annotation within our dataset, particularly in sentences relating to sporting events. Some annotators labeled negative relationships between competing athletes but it seems in sports such relations should not be annotated. Additionally, discrepancies arose among

¹⁴[text-davinci-002-render-paid, 2022/11/22](https://openai.com/research/text-davinci-002-render-paid)

¹⁵<https://pypi.org/project/googletrans/>

annotators concerning the annotation of victories and defeats: while some treated such outcomes as factual information, others interpreted them as positive or negative sentiments towards the athletes. Similar variations were observed in the output from ChatGPT.

9 Conclusion

In this paper, we described the RuSentNE-2023 evaluation devoted to targeted sentiment analysis in Russian news texts. We think that targeted sentiment analysis in news texts has not been thoroughly explored, as prior research primarily focuses on sentiment within user reviews and social media content. The distinct characteristics of news texts include a significant proportion of neutral named entities, a substantial presence of implicit sentiment, and the occurrence of multiple named entities with varying sentiments within the same sentence.

The objective of the RuSentNE-2023 evaluation was to predict sentiment towards a named entity within a single sentence, effectively classifying the sentiment into one of three categories: positive, negative, or neutral. The dataset used for the RuSentNE-2023 evaluation is derived from the Russian News corpus, RuSentNE, which contains a comprehensive sentiment-related annotation. The corpus is marked with named entities and the sentiment directed towards these entities, alongside associated effects and emotional states.

The evaluation was organized using the CodaLab competition framework. The main evaluation measure was macro-averaged measure of positive and negative classes. There were 15 participants in the development stage of the evaluation, with 9 participants presenting their results in the final stage. The best results reached of 66% Macro F-measure for the positive and negative classes.

We additionally conducted experiments with ChatGPT on the test set of our evaluation, discovering that the zero-shot responses of ChatGPT reached 60% of the F-measure, which corresponds to the fourth-place in the evaluation. ChatGPT also provided comprehensive explanations of its conclusions. The outcomes are considerably notable, given that this was a zero-shot application.

Our future plans involve improving the annotation of the RuSentNE corpus, taking into account the results of the participants and the explanations provided by ChatGPT. Subsequently, we aim to conduct an evaluation focused on structured sentiment analysis, specifically tailored for the extraction of four-tuples (the subject of sentiment, object of sentiment, sentiment, and sentiment expression). Furthermore, it is important to eliminate sentence restrictions and advance towards classifying sentiment across sentences.

Acknowledgements

The work was supported by the Russian Science Foundation under Agreement No. 21-71-30003.

References

- Christoph Alt, Marc Hübner, and Leonhard Hennig. 2019a. Improving relation extraction by pre-trained language representations. // *Proceedings of AKBC 2019*.
- Christoph Alt, Marc Hübner, and Leonhard Hennig. 2019b. Improving relation extraction by pre-trained language representations. *arXiv preprint arXiv:1906.03088*.
- Enrique Amigó, Adolfo Corujo, Julio Gonzalo, Edgar Meij, Maarten De Rijke, et al. 2012. Overview of replab 2012: Evaluating online reputation management systems. // *CLEF (online working notes/labs/workshop)*.
- Enrique Amigó, Jorge Carrillo-de Albornoz, Irina Chugur, Adolfo Corujo, Julio Gonzalo, Edgar Meij, Maarten de Rijke, and Damiano Spina. 2014. Overview of replab 2014: author profiling and reputation dimensions for online reputation management. // *Information Access Evaluation. Multilinguality, Multimodality, and Interaction: 5th International Conference of the CLEF Initiative, CLEF 2014, Sheffield, UK, September 15-18, 2014. Proceedings 5*, P 307–322. Springer.
- Jeremy Barnes, Laura Oberländer, Enrica Troiano, Andrey Kutuzov, Jan Buchmann, Rodrigo Agerri, Lilja Øvrelid, and Erik Velldal. 2022. Semeval 2022 task 10: Structured sentiment analysis. // *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, P 1280–1295.

- Mikhail Burtsev, Alexander Seliverstov, Rafael Airapetyan, Mikhail Arkhipov, Dilyara Baymurzina, Nickolay Bushkov, Olga Gureenkova, Taras Khakhulin, Yuri Kuratov, Denis Kuznetsov, Alexey Litinsky, Varvara Logacheva, Alexey Lymar, Valentin Malykh, Maxim Petrov, Vadim Polulyakh, Leonid Pugachev, Alexey Sorokin, Maria Vikhreva, and Marat Zaynutdinov. 2018. DeepPavlov: Open-source library for dialogue systems. // *Proceedings of ACL 2018, System Demonstrations*, P 122–127, Melbourne, Australia, July. Association for Computational Linguistics.
- Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357.
- Iliia Chetviorkin and Natalia Loukachevitch. 2013. Evaluating sentiment analysis systems in russian. // *Proceedings of the 4th biennial international workshop on Balto-Slavic natural language processing*, P 12–17.
- Hyung Won Chung, Thibault Févry, Henry Tsai, Melvin Johnson, and Sebastian Ruder. 2021. Rethinking embedding coupling in pre-trained language models. // *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. // *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, P 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Alvan R Feinstein and Domenic V Cicchetti. 1990. High agreement but low kappa: I. the problems of two paradoxes. *Journal of clinical epidemiology*, 43(6):543–549.
- Anton Golubev and Natalia Loukachevitch. 2020. Improving results on russian sentiment datasets. // *Artificial Intelligence and Natural Language: 9th Conference, AINL 2020, Helsinki, Finland, October 7–9, 2020, Proceedings 9*, P 109–121. Springer.
- Anton Golubev and Natalia Loukachevitch. 2021. Use of augmentation and distant supervision for sentiment analysis in russian. // *Text, Speech, and Dialogue: 24th International Conference, TSD 2021, Olomouc, Czech Republic, September 6–9, 2021, Proceedings 24*, P 184–196. Springer.
- Felix Hamborg, Karsten Donnay, and Bela Gipp. 2021. Towards target-dependent sentiment classification in news articles. // *Diversity, Divergence, Dialogue: 16th International Conference, iConference 2021, Beijing, China, March 17–31, 2021, Proceedings, Part II 16*, P 156–166. Springer.
- Yangkun Lin, Chen Liang, Jing Xu, Chong Yang, and Yongliang Wang. 2022. Zhixiaobao at semeval-2022 task 10: Approaching structured sentiment with graph parsing. // *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, P 1343–1348.
- Natalia Loukachevitch and Anatolii Levchik. 2016. Creating a general russian sentiment lexicon. // *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, P 1171–1176.
- Natalia Loukachevitch and Yuliya Rubtsova. 2015. Entity-oriented sentiment analysis of tweets: results and problems. // *Text, Speech, and Dialogue: 18th International Conference, TSD 2015, Pilsen, Czech Republic, September 14-17, 2015, Proceedings 18*, P 551–559. Springer.
- NV Loukachevitch and N Rusnachenko. 2018. Extracting sentiment attitudes from analytical texts. // *Komp'yuternaja Lingvistika i Intellektual'nye Tehnologii*, P 448–458.
- Natalia Loukachevitch, Pavel Blinov, Evgeny Kotelnikov, Yulia Rubtsova, Vladimir Ivanov, and Elena Tutubalina. 2015. Sentirueval: testing object-oriented sentiment analysis systems in russian. // *Proceedings of International Conference Dialog*, volume 2, P 3–13.
- Natalia Loukachevitch, Ekaterina Artemova, Tatiana Batura, Pavel Braslavski, Iliia Denisov, Vladimir Ivanov, Suresh Manandhar, Alexander Pugachev, and Elena Tutubalina. 2021. Nerel: A russian dataset with nested named entities, relations and events. // *International Conference on Recent Advances in Natural Language Processing: Deep Learning for Natural Language Processing Methods and Applications, RANLP 2021*, P 876–885. Incoma Ltd.
- Dehong Ma, Sujian Li, Xiaodong Zhang, and Houfeng Wang. 2017. Interactive attention networks for aspect-level sentiment classification. // *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI'17*, P 4068–4074. AAAI Press.

- Gaku Morio, Hiroaki Ozaki, Atsuki Yamaguchi, and Yasuhiro Sogawa. 2022. Hitachi at SemEval-2022 task 10: Comparing graph- and Seq2Seq-based models highlights difficulty in structured sentiment analysis. // *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, P 1349–1359, Seattle, United States, July. Association for Computational Linguistics.
- Preslav Nakov, Alan Ritter, Sara Rosenthal, Fabrizio Sebastiani, and Veselin Stoyanov. 2016. Semeval-2016 task 4: Sentiment analysis in twitter. // *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, P 1–18.
- Maria Pontiki, Dimitrios Galanis, Harris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. Semeval-2015 task 12: Aspect based sentiment analysis. // *Proceedings of the 9th international workshop on semantic evaluation (SemEval 2015)*, P 486–495.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammed AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, et al. 2016. Semeval-2016 task 5: Aspect based sentiment analysis. // *ProWorkshop on Semantic Evaluation (SemEval-2016)*, P 19–30. Association for Computational Linguistics.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. Semeval-2017 task 4: Sentiment analysis in twitter. // *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, P 502–518.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts. // *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, P 4222–4235, Online, November. Association for Computational Linguistics.
- Chi Sun, Luyao Huang, and Xipeng Qiu. 2019a. Utilizing BERT for aspect-based sentiment analysis via constructing auxiliary sentence. // *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, P 380–385, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Chi Sun, Luyao Huang, and Xipeng Qiu. 2019b. Utilizing bert for aspect-based sentiment analysis via constructing auxiliary sentence. // *Proceedings of NAACL-HLT*, P 380–385.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Biqing Zeng, Heng Yang, Ruyang Xu, Wu Zhou, and Xuli Han. 2019. Lcf: A local context focus mechanism for aspect-based sentiment classification. *Applied Sciences*, 9(16):3389.
- Bowen Zhang, Daijun Ding, and Liwen Jing. 2022. How would stance detection techniques evolve after the launch of chatgpt? *arXiv preprint arXiv:2212.14548*.
- Liu Zhuang, Lin Wayne, Shi Ya, and Zhao Jun. 2021. A robustly optimized BERT pre-training approach with post-training. // *Proceedings of the 20th Chinese National Conference on Computational Linguistics*, P 1218–1227, Huhhot, China, August. Chinese Information Processing Society of China.