



# Automatic Sentiment Analysis of Texts: The Case of Russian

*Natalia Loukachevitch*

## 28.1 INTRODUCTION

Automatic sentiment analysis of texts, that is, the identification of the author's opinion about the subject discussed in the text, has been one of the most significant tasks in natural language processing in the past two decades. The interest in sentiment analysis is connected to the large volume of electronic texts available on social networks and online recommendation services that contain an abundance of individuals' opinions on various issues: from products and services to the current political and economic situation (for more on corpora, see Chap. 17).

A large number of scholarly works is devoted to sentiment analysis of user reviews stored in recommendation services (Pang et al. 2002; Pang and Lee 2008; Liu 2012). Another important area of sentiment analysis is the so-called reputation monitoring that tracks positive and negative feedback about a company and its products (Amigo et al. 2012). Sentiment analysis of financial reports and financial news is used to determine trends in the stock and currency markets (Nassirtoussi et al. 2015). The sentiment of mentioning terms in scientific articles is used to predict the most important concepts and scientific trends (McKeown et al. 2016). Sentiment information extracted from texts can be used to determine the personal characteristics of the author (Volkova et al. 2015).

The role of automatic sentiment analysis of social network messages for political and social research is growing. Such studies include the identification

---

N. Loukachevitch (✉)

Lomonosov Moscow State University, Moscow, Russia

© The Author(s) 2021

D. Gritsenko et al. (eds.), *The Palgrave Handbook of Digital Russia Studies*, [https://doi.org/10.1007/978-3-030-42855-6\\_28](https://doi.org/10.1007/978-3-030-42855-6_28)

501

of political preferences (Volkova et al. 2014), the prediction of election results (Vepsäläinen et al. 2017; Vilares et al. 2015), and the identification of attitudes toward various political decisions. Also, automatic sentiment analysis can be used to recognize hate speech and calls for violence or fake news (Volkova and Bell 2016).

The first approaches to sentiment analysis aimed to determine the overall sentiment of the document or its fragment (Pang et al. 2002). This level of analysis assumes that a document expresses a unanimous opinion about a single entity, such as in a review of a product. Since the document can express multiple attitudes in relation to the different entities it contains, at the next stage scholars studied the tasks of sentiment analysis aimed toward specified entities mentioned in the text (Amigo et al. 2012; Jiang et al. 2011; Loukachevitch et al. 2015; Loukachevitch and Rubtsova 2016). Finally, an even more detailed level of sentiment analysis is the analysis of opinions on specific properties or parts (the so-called aspects) of the entity (Liu and Zhang 2012; Pontiki et al. 2016; Popescu and Etzioni 2007).

Liu and Zhang (2012, 4) define opinion as a five tuple  $(e_i, a_{ij}, s_{ijkl}, h_k, t_l)$ , where  $e_i$  is the name of an entity to which the opinion relates,  $a_{ij}$  is an aspect (part or characteristic) of  $e_i$ ,  $s_{ijkl}$  is the sentiment regarding the entity and its aspect,  $h_k$  is the author of the opinion (opinion holder), and  $t_l$  is the time when the opinion is expressed by  $h_k$ . The sentiment  $s_{ijkl}$  may be positive, negative, or neutral, or may be expressed with varying degrees of intensity that is measured, for example on a scale of 1–5.

In this chapter, we first describe the problems that can be encountered in automatic sentiment analysis. Then, we briefly consider the main methods to conduct sentiment analysis and approaches to creating sentiment vocabularies. Finally, Russian-specific components of automatic sentiment analysis are described, including publicly available vocabularies and sentiment-related shared tasks.

## 28.2 PROBLEMS IN SENTIMENT ANALYSIS

If we ask native speakers what the most significant problems in sentiment analysis would be, the respondents often name irony and sarcasm. Certainly, the difficulties with these language phenomena really exist but problems of automatic sentiment analysis are much more diverse. In what follows, six additional challenges of sentiment analysis are presented.

### 28.2.1 *Multiple Opinions in a Single Text*

Approaches to extracting the main components of opinion largely depend on the genre of the analyzed text. One of the most studied genres of text in the task of sentiment analysis are user reviews on products or services. Such texts usually consider a single entity (but, perhaps in its different aspects), and the

opinion is expressed by one author, namely the reviewer (Pang et al. 2002; Pang and Lee 2008; Liu 2012).

Another popular type of texts for sentiment extraction is Twitter messages (Pak and Paroubek 2010; Rosenthal et al. 2017; Loukachevitch and Rubtsova 2016). Tweets (Twitter posts) were limited to 140 symbols before 2017, when they were extended to 280 characters. Such short texts often require precise sentiment analysis but most of them mention the only opinion target and opinion holder (for more on Twitter analysis, see Chap. 30). The following tweet shows an example of a negative attitude toward Russian phone company Megafon, presented in sarcastic form, which requires the use of sophisticated methods to reveal the correct attitude:

*Megafon, spasibo tebe za zablockirovannye uvedomleniâ ot Rajffajzena* [Megafon, thank you for the blocked notifications from Raiffeisen]

It can seem that in longer texts the author's opinion can be repeated several times in different ways, which would facilitate the analysis. However, long texts may include various entities and related sentiments (Choi et al. 2016; Loukachevitch and Rusnachenko 2018) and they may mention opinions of different persons. If the task is to find an attitude toward the entities mentioned, then the problem of determining the scope of the sentiments arises. For example, sentiment extraction is often carried out in relation to an entity mentioned in the same sentence. However, the author can refer to an entity using the means of reference, for example, pronouns. In addition, if the entire text is devoted to the discussion of one entity, then it can be explicitly mentioned far from the sentiment location (Ben-Ami et al. 2014).

In such document genres as news texts, or especially analytical texts, many opinions from different sources can be simultaneously mentioned. These texts contain opinions conveyed by different subjects, including the author(s)' attitudes, the positions of cited sources, and the relations of the mentioned entities to each other. Analytical texts usually contain a lot of named entities, and only a few of them are subjects or objects of a sentiment attitude (Loukachevitch and Rusnachenko 2018). It is clear that in texts with multiple subjects and/or objects of opinion, the complexity of high-quality automatic analysis of sentiment increases manifold.

### 28.2.2 *Implicit vs. Explicit Sentiment*

It is usually assumed that sentiment is expressed using specialized sentiment words (such as *good*, *bad*, *awful*), which is an explicit way of conveying attitudes. However, sentiment can be expressed also implicitly with the so-called sentiment facts (Liu 2012; Loukachevitch and Levchik 2016; Tutubalina 2015) or words with connotations (Feng et al. 2013).

According to the definition provided by Liu (2012, 26), an implicit opinion is an objective statement, from which the sentiment follows, that is, an implicit

opinion that conveys a desirable or undesirable fact. In preparation of datasets for testing sentiment analysis systems, such sentiment facts can be specifically annotated (Loukachevitch et al. 2015; Nozza et al. 2017). For example, Russian restaurant reviews may include such sentences as: “*Dolgo žдали*” (Waited for a long time) or “*Našli muhu v supe*” (Found a fly in the soup), which, on the one hand, describe what happened (report real facts), but on the other hand convey sentiment.

Connotation is a feeling or idea that is suggested by a particular word, although it need not be a part of the word’s meaning. Connotations often convey positive or negative sentiment (Feng et al. 2013). The appearance of words with positive or negative connotations in a text correlates with the corresponding sentiment expressed in the text. For example, in movie reviews, names of famous actors usually have positive connotations. In restaurant reviews, the noun *muha* (fly) is associated with a negative sentiment in different contexts, for example:

*No sil’no dulo ot okna, pri ètom letala nazojlivaâ muha i ne hvatalo oficiantov* [But there was a strong draft from the window, while an annoying fly was flying around and there were not enough waiters].

*Prišli v kafe na Ozernoj, oficiantku ele doždalis’ ležala muha mertvaâ na stole* [Went to the cafe on Lake street, barely waited for the waitress, there was a dead fly on the table].

An interesting example of a word with specific connotations in Russian restaurant reviews is the word *majonez* (mayonnaise). Many sources indicate *majonez* as a key component of Soviet and Russian cuisine (Shearlaw 2014; Whalley 2018). However, when mentioned in contemporary Russian restaurant reviews, this word usually conveys negative sentiment, for example:

*Absolútno vse salaty soderžat majonez, pričem ego vezde mnogo* [Absolutely all salads contain mayonnaise, and lots of it in everything].

*Edinstvennye teplye rolly byli táželovaty vvidu nalíciâ v nih majoneza* [The only hot rolls were heavy due to the presence of mayonnaise].

In news and analytical texts, we can find a lot of words with international negative connotations such as *war*, *unemployment*, *segregation*, or *traffic jam*. Positive connotations are often associated with achievements of a nation. For example, in Russia positive connotations are associated with cosmos-related concepts such as *sputnik*, *Yuri Gagarin*, or *MKS* (International Space Station).

Gradual adjectives (such as *long – short*, *large – small*, etc.) can often convey sentiment facts but their sentiment orientation is very dependent of the context (Cambria et al. 2010). For example, the word *long* can be both negative and positive in the digital camera domain: if it has a long battery life, it means the battery is good; if you need to adjust the focus for a long time, then the opinion about the camera is negative.

Because of the existence of implicit sentiments and connotations, it is impossible to create general sentiment lexicons, which can be equally useful across many domains. It is, therefore, necessary to develop specialized sentiment lexicons using domain-specific text collections or update existing general lexicons to adapt them to a specific domain. (Hamilton et al. 2016; Severyn and Moschitti 2015; Chetviorkin and Loukachevitch 2012).

### 28.2.3 *Ambiguity of Sentiment Words*

Difficulties with the interpretation of explicit sentiment vocabulary may also arise. Sentiment words can be ambiguous: in one sense, they can be neutral, while in other senses they are negative or positive (Akkaya et al. 2009; Baccianella et al. 2010). For example, the Russian word *presnyj* (fresh) bears a positive connotation in the phrase *presnaá voda* (freshwater), while in other senses of the word *presnyj* (tasteless for food and uninteresting as in movie reviews) this word is negative.

A word can change or lose its polarity depending on the subject area or the current context. For example, the Russian sentiment words *verolomstvo* (treachery) and *predatel'stvo* (betrayal) cannot be considered as conveying an opinion in movie reviews, because they are usually mentioned in a movie synopsis to retell the plot of a movie. The word *smešnoj* (funny), most likely, is negative in the sphere of politics, yet indicates a positive orientation when it is used in reviews of comedies. When characterizing other movie genres, this word can be both positive and negative.

### 28.2.4 *Sentiment Modifiers*

The appearance of sentiment words in the text may be accompanied by sentiment modifiers that enhance (for example, *much*, *more*), reduce (*too*, *less*) or inverse prior word sentiment (negation: *no*, *not*). Thus, when analyzing the sentiment, such modifiers should be taken into account, and it is necessary to have some numerical model that modifies the original polarities of the word (Taboada et al. 2011; Wilson et al. 2005; Wiegand et al. 2010). One of the common models of accounting for polarity modifiers ascribes some coefficients to them, which are considered as factors modifying the initial polarity of the words to which these modifiers relate.

Another important issue is determining the scope of the polarity modifier in a particular sentence (Taboada et al. 2011). Most approaches suppose that polarity modifiers, such as negation, modify sentiment of neighbor words, but long-distance influence is also possible. For example, in the sentence “*Á ne dumaú, čto èto zasluživayet upominaniá*” (I do not think it is worth mentioning), the negation changes the sentiment orientation of the word *zasluživayet* (worth) from positive to negative.

If negation stands before several sentiment-bearing words, it is important to calculate the overall sentiment of the whole group and then to apply negation

to it. In the following sentence, we see the phrase “*ne boitsá raskola*” (is not afraid of a split), where negation stands before two words with negative sentiment. To obtain the positive mood of the sentence, it is necessary to determine the sentiment of the phrase as negative and then to apply negation to it:

*Sekretar’ prezidiuma gensoвета “Edinoj Rossii,” zampredsedatelá Gosdumy Sergej Neverov v subbotu zaávil, čto partiá ne boitsá raskola v svázi s poávleniem v nej raznyh ideologičeskikh platform* [Secretary of the Presidium of the General Council of “United Russia,” Deputy Chairman of the State Duma Sergei Neverov, on Saturday stated that the party **is not afraid of a split** in connection to the appearance of various ideological platforms within it].

Polarity modifiers can also form groups such as double negation. We can see such double negation *ne bez* in well-known Russian proverb “*V sem’ye ne bez uroda,*” which translates into English without any negations: “Every family has its black sheep.” In this example, we see negative sentiment as if negation coefficients were multiplied.

### 28.2.5 Factors of Irreal Context

When analyzing the sentiment, it is important to consider how a proposition conveying sentiment corresponds to reality (Saurí and Pustejovsky 2012; Taboada et al. 2011; Wilson et al. 2005). For example, in the sentence “*My nadeáliš’, čto nam ponravitsá kino*” (We hoped that we would like the movie), one can see the positive word *ponravitsá* (like), but it says nothing about whether the author really liked the movie.

In linguistics, this is covered by the concept of irreals or irreal mood, which is a group of grammatical means that is used to denote that what is said in a sentence does not refer to what really happens (Taboada et al. 2011). In every language, there are some factors showing that the proposition is not factual (the so-called irrealis markers). In Russian modal verbs, private-state verbs, such as *nadeát’sá* (to hope), *ožídat’* (to expect), *dumat’* (to think), can be used as such markers.

According to Kuznetsova et al. (2013, 72), for the Russian language, such function words as *esli*, *by*, *li*, *esli by* also often mark the irrealis mood. When selecting parameters on the training set, Kuznetsova et al. (2013, 72) indicated that the prior sentiment scores of sentiment words found in the sentences with irrealis markers should be decreased (but not nullified).

### 28.2.6 Comparisons

Comparisons complicate the process of determining sentiment, because additional entities are mentioned in the text, and some sentiments can refer to them. It is often supposed that comparisons are conveyed with the so-called comparative constructions such as *lučše čem* (better than) or *dorože čem* (more

expensive than). In most cases, comparisons may be introduced without any specialized constructions. Additional entities mentioned for comparison are sometimes very difficult to detect, and it can also be a complex task to single out the attitudes related to them. For example, in the following extract from a restaurant review, the comparison is marked with word *drugoy* (another), and positive words *naslaždalis'* (enjoyed) and *volšebnym* (wonderful) characterize a restaurant distinct from the restaurant under review (example from Loukachevitch et al. 2015, 8):

*My rešili ne brat' zdes' desert i kofe, a pošli v drugoj restoran, gde naslaždalis' volšebnym zaveršeniem našego večera* (We decided not to have dessert and coffee there, but instead went to another restaurant where we **enjoyed a wonderful** end to our evening).

### 28.2.7 Irony and Sarcasm

The processing of irony and sarcasm is a serious problem for sentiment analysis systems, since the sentiment of an ironic (sarcastic) utterance differs from its literal sentiment (Wilson and Sperber 2007). In Benamara et al. (2017, 37), a generalized understanding of irony is proposed as “an incongruity between the literal meaning of an utterance and its intended meaning.” Most often, a positive-looking statement (containing more positive sentiment words or an equal number of positive and negative words) hides a negative opinion, for example, “*Sberbank—naibolee krupnaâ set' nerabotaûsib bankomatov v Rossii*” (Sberbank is the largest network of nonoperating ATMs in Russia). Sarcasm is regarded as a sharper, more aggressive, possibly degrading form of irony (Benamara et al. 2017).

The annotation of textual data for the study of irony and sarcasm is a complex task. Interesting data for analyzing these phenomena are Twitter messages that the user can mark with special hashtags: *#irony*, *#sarcasm* and some others (Reyes et al. 2013; Sulis et al. 2016). However, recent studies of irony in Twitter show that ironic tweets marked with hashtags and annotated by experts have different characteristics (Kunneman et al. 2015). In addition, in the Russian segment of Twitter, users do not use similar Russian hashtags in the same way as American or European audience (Zefirova and Loukachevitch 2019, 48). The “*ironiâ*” (irony) hashtag is mostly used as a description for images or jokes alongside with such hashtags as “*#šutka*” (joke), “*#smel*” (laugh) and does not seem to express the desired content.

## 28.3 METHODS AND RESOURCES USED IN SENTIMENT ANALYSIS

Automatic analysis of sentiment can utilize two main types of approaches (Liu 2012; Pang and Lee 2008): knowledge-based methods using sentiment lexicons and rules (Taboada et al. 2011; Kuznetsova et al. 2013) and approaches

based on machine learning (Liu 2012; Pang and Lee 2008). Knowledge-based methods require the creation of a specialized sentiment lexicon for a specific domain. Linguistic rules are necessary to sum up sentiment scores of several sentiment words and for accounting for the word context (sentiment modifiers, unreal context, etc.).

Supervised machine learning requires preliminary annotation of a training collection. Depending on the task, different classification algorithms, features of the text representation, and feature weights can be chosen (Pang et al. 2002; Pang and Lee 2008; Liu 2012). Currently, the best results in machine-learning sentiment analysis are achieved by deep learning with neural networks (Rosenthal et al. 2017; Cliché 2017; Arkhipenko et al. 2016), which substituted a previous leader: Support vector machine (SVM) classifier (Pang et al. 2002; Pang and Lee 2008; Chetviorkin and Loukachevitch 2013).

At present, there exist also approaches that integrate available sentiment vocabularies (both manually created and automatically generated) into machine learning methods, transforming them into specialized features (Rosenthal et al. 2017; Mohammad et al. 2013; Loukachevitch and Levchik 2016). The use of preliminary created lexicons helps to overcome data sparsity of training collections (Loukachevitch and Rubtsova 2016). Below, we consider some approaches to creating sentiment lexicons and publicly available Russian lexicons.

Most sentiment vocabularies look like lists of words and expressions with scores of their sentiment (Wilson et al. 2005). Some vocabularies also provide additional characteristics of the word sentiment called “strength.” Sentiment scores can also be assigned to specific senses of ambiguous words (Baccianella et al. 2010; Loukachevitch and Levchik 2016).

For many languages, general sentiment vocabularies have been published. Despite the fact that in each particular domain, specialized vocabularies are needed, general lexicons are also useful since they can serve as source material, which can then be adapted to a domain. Domain-specific sentiment vocabularies are usually generated with automatic or semiautomatic methods using domain-specific text collections (Hamilton et al. 2016; Severyn and Moschitti 2015; Chetviorkin and Loukachevitch 2012).

For Russian, Chetviorkin and Loukachevitch (2012) have described an automatically generated Russian sentiment lexicon in the domain of products and services called ProductSentiRus (ProductSentiRus 2012). The ProductSentiRus lexicon is obtained by applying a supervised machine-learning model to several domain review collections. It is presented as a list of 5000 words ordered by the decreased probability of their sentiment orientation without any positive or negative labels. For example, the most probable sentiment words in ProductSentiRus are as follows: *bespodobnyj* (peerless), *nevñatnyj* (slurred), *obaldennyj* (awesome), *otvratnyj* (disgusting), et cetera.

The general Russian lexicon of sentiment words and expressions, RuSentiLex (RuSentiLex 2017), was created in a semiautomatic way (Loukachevitch and Levchik 2016). The entries of the RuSentiLex lexicon are classified according to four sentiment categories (positive, negative, neutral, or positive/negative)



and three sources of sentiment (opinion, emotion, or fact). The words in the lexicon that have different sentiment scores in different senses are linked to the appropriate concepts of the thesaurus of the Russian-language RuThes (Loukachevitch and Dobrov 2014; RuThes 2016), which can help disambiguate sentiment ambiguity in specific domains or contexts. The lexicon was gathered from several sources: opinionated words from general Russian thesaurus RuThes, slang and curse words extracted from Twitter, and objective words with positive or negative connotations from a news collection (Loukachevitch and Levchik 2016; for more on RuThes, see Chap. 18). For example, the description of word *presnyj* in RuSentiLex is as follows (labels in quotes correspond to the names of RuThes concepts):

- *presnyj*, Adj, *presnyj*, negative, emotion, “NEVKUSNYJ” [*tasteless*];
- *presnyj*, Adj, *presnyj*, negative, opinion, “NEINTERESNYJ” [*insipid*];
- *presnyj*, Adj, *presnyj*, positive, fact, “PRESNAĀ VODA” [*fresh water*]

The Russian sentiment lexicon LINIS Crowd was created via crowdsourcing (Koltsova et al. 2016; LINIS Crowd SENT 2016). The lexicon is aimed at detecting sentiment in user-generated content (blogs, social media) related to social and political issues. Each word was assessed by at least three volunteers in the context of different texts. The words were scored from  $-2$  (negative) to  $+2$  (positive). For example, the word *anarhizm* (anarchism) obtained three 0 (neutral) scores and three  $-1$  (weakly negative) scores in the considered contexts.

Several international lexicons were automatically constructed for Russian. The Chen-Skienna’s lexicon (2876 words) (Chen and Skienna 2014; Chen-Skienna’s Lexicon 2014) was generated for 136 languages via graph propagation from seed words. However, from the human point of view, the words included in this automatically generated lexicon seem extremely strange. For example, positive words in the Chen-Skienna’s Lexicon include such words as *tipa* (type of), *post* (post), *sootvetstvenno* (correspondingly), *sovsem* (at all), et cetera.

Mohammad and Turney (2013) generated the Russian variant of the EmoLex lexicon (EmoLex 2017) with automatic translation from the English lexicon obtained by crowdsourcing (4412 Russian words).

Kotelnikov et al. (2018) studied available Russian sentiment lexicons and found that all the lexicons have relatively small intersection with each other. Besides, the translated lexicons (EmoLex and Chen-Skienna’s lexicon) have a smaller intersection with other lexicons than on average (10.0%), and at the same time are relatively similar to each other (18.2%). Kotelnikov et al. (2018) also compared available Russian lexicons as features in machine-learning text categorization. Users’ reviews from five domains (books, movies, banks, hotels, and kitchens) were used as text collections for the experiments. The study found that the best results of classification using a single lexicon in all domains

were obtained with ProductSentiRus (Chetviorkin and Loukachevitch 2012). The union of all lexicons gives slightly better results.

As was mentioned, useful sentiment lexicons should be fine-tuned or constructed specially for the domain under analysis. Therefore, to apply sentiment lexicons in a specific domain, it is recommended to gather all available lexicons and to collect a domain-specific text collection as large as possible. Having such data, it is possible to filter out sentiment words and constructions relevant in the domain.

## 28.4 RUSSIAN SENTIMENT-RELATED SHARED TASKS

For the evaluation of Russian sentiment analysis systems, several shared tasks have been organized. In 2011–2013, two evaluations of document-level sentiment approaches were carried out. Two types of text collections were used for the evaluation: users' reviews in three domains (movies, books, and digital cameras) and news quotations (Chetviorkin and Loukachevitch 2013).

For training in the review track, users' reviews were collected from recommendation services (Imhonet and Yandex.market). The reviews had users' scores on a ten-point scale for the Imhonet reviews (movies and books) and on a five-point scale for the Yandex reviews (digital cameras). The participants could choose any of tracks classifying reviews into two, three, or five classes. The reviews for the test collections were extracted from social network messages. The sentiment annotation was created manually by human experts. The participants utilized various machine-learning and knowledge-based approaches, but the best methods in all review-related tasks were SVM-based classifiers.

In the quotation track, direct or indirect speech fragments extracted from news reports had to be classified into three classes (positive, negative, or neutral). About 5000 fragments each were prepared for the training and test collection. Both collections were annotated manually; therefore, the size of the training collection was much smaller than for the review task. In this quotation task, the knowledge-based approaches showed the best results (Chetviorkin and Loukachevitch 2013).

The second series of Russian sentiment analysis evaluations (SentiRuEval 2014–2016) was devoted to the entity-oriented and aspect-based tasks of sentiment analysis. Namely, the tasks included aspect-based analysis of reviews in two domains (car and restaurant reviews). Using the prepared collections, Russian training and test datasets were further utilized in the international SemEval aspect-based sentiment evaluation in 2016 (Pontiki et al. 2016; ABSA SemEval-2006 2016).

The entity-oriented task was based on Twitter messages. The participants were asked to classify messages into three classes from the point of view of reputation monitoring (positive, negative, or neutral) in two separate domains: banks and mobile operators. For example, positive tweets could contain a

positive opinion or positive fact about the company. The training and test collections were prepared via crowdsourcing (SentiRuEval-2016 data 2016).

The approaches of the participants for the Twitter sentiment analysis differed significantly in 2015 and 2016 (Loukachevitch and Rubtsova 2016). In 2015, the basic approach was the SVM classifier trained on only the training collection without any additional data (unlabeled text collections or sentiment lexicons). Due to this, the participating systems could make mistakes in the classification of the test tweets if a tweet contained sentiment words absent in the training dataset (Loukachevitch and Rubtsova 2016).

In 2016, the best approach was based on neural networks, which used word embeddings (vector representations of words) calculated on a large collection of user comments (Arkhipenko et al. 2016). Such representations allowed the winner to overcome the differences in the training and test collections because words that have semantic similarity also have similar vector representations. The next most successful approaches in terms of the quality of results combined machine learning and the existing Russian lexicons (Loukachevitch and Rubtsova 2016).

## 28.5 CONCLUSION

Automatic sentiment analysis of texts is among the popular applications in natural language processing of texts. In this chapter, we described the problems that can be encountered in automatic sentiment analysis. Then, we briefly considered the main methods for sentiment analysis and approaches to creating sentiment vocabularies. Finally, Russian-specific components of automatic sentiment analysis—publicly available vocabularies and sentiment-related shared tasks—were presented.

The current state of affairs in sentiment analysis (including in its application to the Russian language) can be characterized as follows: approaches to sentiment analysis of some text genres, such as user reviews or short posts on social networking sites, are well studied, but there are a lot of complicated phenomena in sentiment analysis that require further research, especially in the processing of full-text news and analytical articles.

From the practical point of view, there are at least four Russian sentiment vocabularies currently available on the Internet. To apply sentiment lexicons in a specific domain, it is recommended to gather all available lexicons and to collect a domain-specific text collection as large as possible. Having such data, it is possible to filter out sentiment words and constructions that are relevant in the domain.

## REFERENCES

- ABSA SemEval-2016. 2016. Data for Aspect-Based Sentiment Analysis, SemEval-2016. <http://alt.qcri.org/semeval2016/task5/index.php?id=data-and-tools>.

- Akkaya, Cem, Janyce Wiebe, and Rada Mihalcea. 2009. Subjectivity Word Sense Disambiguation. *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, vol. 1, 190–199. Association for Computational Linguistics.
- Amigo, Enrique, Adolfo Corujo, Julio Gonzalo, Edgar Meij, and Maarten Rijke. 2012. Overview of RepLab. 2012: Evaluating Online Reputation Management Systems. *CLEF-2012 Working Notes*. <http://ceur-ws.org/Vol-1178/CLEF2012wn-RepLab-AmigoEt2012.pdf>.
- Arhipenko, Konstantin, Ilya Kozlov, Yuriy Trofimovich, Kirill Skorniakov, Andrey Gomzin, and Denis Turdakov. 2016. Comparison of Neural Network Architectures for Sentiment Analysis of Russian Tweets. *Proceedings of International Conference on computational linguistics and intellectual technologies Dialog-2016*, 50–58.
- Baccianella, Stefano, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. *Proceedings of Language Resources and Evaluation Conference LREC-2010*, vol. 10, 2200–2204.
- Benamara, Farah, Maite Taboada, and Yannick Mathieu. 2017. Evaluative Language Beyond Bags of Words: Linguistic Insights and Computational Applications. *Computational Linguistics* 43: 201–264.
- Ben-Ami, Zvi, Ronen Feldman, and Binyamin Rosenfeld. 2014. Entities' Sentiment Relevance. *Proceedings of Association for Computational Linguistics Conference ACL-2014*, 87–92.
- Cambria, Eric, Amir Hussain, Catherine Havasi, and Chris Eckl. 2010. Sentic Computing: Exploitation of Common Sense for the Development of Emotion-Sensitive Systems. *Development of Multimodal Interfaces: Active Listening and Synchrony* 5967: 148–156. Berlin and Heidelberg: Springer, LNCS.
- Chen, Yanqing, and Steven Skiena. 2014. Building Sentiment Lexicons for All Major Languages. *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics ACL-2014*, vol. 2, 383–389.
- Chen-Skienna's Lexicon. 2014. Multilingual Sentiment Lexicons, Including Russian. <https://sites.google.com/site/datascienceslab/projects/multilingualsentiment>.
- Chetviorkin, Iliia, and Natalia Loukachevitch. 2012. Extraction of Russian Sentiment Lexicon for Product Meta-Domain. *Proceedings of COLING-2012*, 593–610.
- . 2013. Evaluating Sentiment Analysis Systems in Russian. *Proceedings of the 4th Biennial International Workshop on Balto-Slavic natural Language Processing*, 12–17.
- Choi, Eunsol, Hannah Rashkin, Luke Zettlemoyer, and Yejin Choi. 2016. Document-level Sentiment Inference with Social, Faction, and Discourse Context. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL-2016*, 333–343.
- Cliché, Mathieu. 2017. BB twtr at SemEval-2017 Task 4: Twitter Sentiment Analysis with CNNs and LSTMs. *Proceedings of the 11th International Workshop on Semantic Evaluation SemEval 17*, 572–579.
- EmoLex. 2017. NRC Word-Emotion Association Lexicon, Version 2017. <http://www.saifmohammad.com/WebPages/NRC-Emotion-Lexicon.htm>.
- Feng, Song, Jun Seok Kang, Polina Kuznetsova, and Yejin Choi. 2013. Connotation Lexicon: A Dash of Sentiment Beneath the Surface Meaning. *Proceedings of the 51th Annual Meeting of the Association for Computational Linguistics, ACL-2013*, 1774–1784.
- Hamilton, William, Kevin Clark, Jure Leskovec, and Dan Jurafsky. 2016. Inducing Domain-Specific Sentiment Lexicons from Unlabeled Corpora. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 595–604.

- Jiang, Long, Mo Yu, Ming Zhou, Xiaohua Liu, and Tiejun Zhao. 2011. Target Dependent Twitter Sentiment Classification. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics ACL-2011*, 151–160.
- Koltsova, Olesya, Svetlana Alexeeva, and Sergey Kolcov. 2016. An Opinion Word Lexicon and a Training Dataset for Russian Sentiment Analysis of Social Media. *Proceedings of Computational Linguistics and Intellectual Technologies Conference Dialogue-2016*, 277–287.
- Kotelnikov, Evgeny, Tatiana Peskischeva, Anastasia Kotelnikova, and Elena Razova. 2018. A Comparative Study of Publicly Available Russian Sentiment Lexicons. *Conference on Artificial Intelligence and Natural Language. AINL 2018. Communications in Computer and Information Science 930*, 139–151. Cham: Springer.
- Kunneman, Florian, Christine Liebrecht, Margot van Mulken, and Antal van den Bosch. 2015. Signaling Sarcasm: From Hyperbole to Hashtag. *Information Processing and Management* 51: 500–509.
- Kuznetsova, Ekaterina, Natalia Loukachevitch, and Ilya Chetviorkin. 2013. Testing Rules for a Sentiment Analysis System. *Proceedings of International Conference on Computational Linguistics and Intellectual Technologies Dialog-2013*, vol. 2, 71–81.
- LINIS crowd SENT. 2016. Russian Sentiment Lexicon, Version of 2016. <http://linis-crowd.org/>.
- Liu, Bing. 2012. *Sentiment Analysis and Opinion Mining*. Morgan & Claypool Publishers.
- Liu, Bing, and Lei Zhang. 2012. A Survey of Opinion Mining and Sentiment Analysis. In *Mining Text Data*, 415–463. Springer.
- Loukachevitch, Natalia, and Boris Dobrov. 2014. RuThes Linguistic Ontology vs. Russian Wordnets. *Proceedings of the Seventh Global Wordnet Conference GWC-2014*, 154–162.
- Loukachevitch, Natalia, and Anatoly Levchik. 2016. Creating a General Russian Sentiment Lexicon. *Proceedings of Language Resources and Evaluation Conference LREC-2016*, 1171–1176.
- Loukachevitch, Natalia, and Yuliya Rubtsova. 2016. SentiRuEval-2016: Overcoming Time Gap and Data Sparsity in Tweet Sentiment Analysis. *Proceedings of the Annual International Conference Dialogue-2016*, 416–427.
- Loukachevitch, Natalia, and Nicolay Rusnachenko. 2018. Extracting Sentiment Attitudes from Analytical Texts. *Proceedings of Computational Linguistics and Intellectual Technologies, Papers from the Annual Conference Dialog-2018*, 459–468.
- Loukachevitch, Natalia, Pavel Blinov, Evgeny Kotelnikov, Yuliya Rubtsova, Vladimir Ivanov, and Elena Tutubalina. 2015. SentiRuEval: Testing Object-Oriented Sentiment Analysis Systems in Russian. *Proceedings of International Conference of Computational Linguistics and Intellectual Technologies Dialog-2015*, vol. 2, 2–13.
- McKeown, Kathy, Hal Daume, Snigdha Chaturvedi, John Paparrizos, Kapil Thadani, Pablo Barrio, and Luis Gravano. 2016. Predicting the Impact of Scientific Concepts Using Full Text Features. *Journal of the Association for Information Science and Technology* 67 (11): 2684–2696.
- Mohammad, Saif, and Peter D. Turney. 2013. Crowdsourcing a Word-Emotion Association Lexicon. *Computational Intelligence* 29 (3): 436–465.
- Mohammad, Saif, Svetlana Kiritchenko, and Xiaodan Zhu. 2013. NrcCanada: Building the State-of-the-Art in Sentiment Analysis of Tweets. *Proceedings of Second Joint Conference on Lexical and Computational Semantics (\*SEM)*, vol. 2, 321–327.

- Nassirtoussi, Arman K., Saeed Aghabozorgi, Teh YingWah, and David Ngo. 2015. Text Mining of News-Headlines for FOREX Market Prediction: A Multi-layer Dimension Reduction Algorithm with Semantics and Sentiment. *Expert Systems with Applications* 42 (1): 306–324.
- Nozza, Debora, Elisabetta Fersini, and Enza Messina. 2017. A Multi-view Sentiment Corpus. *Proceedings of EACL-2017*, 273–280.
- Pak, Alexander, and Patrick Paroubek. 2010. Twitter as a Corpus for Sentiment Analysis and Opinion Mining. *Proceedings of Language Resources and Evaluation Conference LREC-2010*, 1320–1326.
- Pang, Bo, and Lillian Lee. 2008. Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval* 2 (1–2): 1–135.
- Pang, Bo, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs Up?: Sentiment Classification Using Machine Learning Techniques. *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing*, vol. 10, 79–86.
- Pontiki, Maria, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammad AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphee De Clercq, Veronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia Loukachevitch, Evgeniy Kotelnikov, Núria Bel, Salud María Jiménez-Zafra and Gülşen Eryiğit. 2016. SemEval-2016 Task 5: Aspect Based Sentiment Analysis. *Proceedings of the 10th International workshop on Semantic Evaluation, Semeval-2016*, 19–30.
- Popescu, Ana-Maria, and Orena Etzioni. 2007. Extracting Product Features and Opinions from Reviews. In *Natural Language Processing and Text Mining*, 9–28. London: Springer.
- ProductSentiRus. 2012. Russian Sentiment Lexicon for Product and Services. <http://www.labinform.ru/pub/productsentirus/productsentirus.txt>.
- Reyes, Antonio, Paolo Rosso, and Tony Veale. 2013. A Multidimensional Approach for Detecting Irony in Twitter. *Language Resources and Evaluation* 47: 1–30.
- Rosenthal, Sara, Noara Farra, and Preslav Nakov. 2017. SemEval-2017 Task 4: Sentiment Analysis in Twitter. *Proceedings of the 11th International Workshop on Semantic Evaluation SemEval-2017*, 502–518.
- RuSentiLex. 2017. Russian Sentiment Lexicon, Version of 2017. [http://www.labinform.ru/pub/rusentilex/rusentilex\\_2017.txt](http://www.labinform.ru/pub/rusentilex/rusentilex_2017.txt).
- RuThes. 2016. Thesaurus of Russian Language, Version of 2016. [http://www.labinform.ru/pub/ruthes/index\\_eng.htm](http://www.labinform.ru/pub/ruthes/index_eng.htm).
- Saurí, Roser, and James Pustejovsky. 2012. Are you sure that this happened? Assessing the Factuality Degree of Events in Text. *Computational Linguistics* 38 (2): 261–299.
- SentiRuEval-2016 data. 2016. Training and Test Collections for Tweet Classification in Russian. <https://goo.gl/GhX3vU>.
- Severyn, Aliaksei, and Alessandro Moschitti. 2015. On the Automatic Learning of Sentiment Lexicons. *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics*, 1397–1402.
- Shearlaw, Maeve. 2014. Understanding Russia's Obsession with Mayonnaise. *The Guardian*. <https://www.theguardian.com/world/2014/nov/21/-sp-understanding-russias-obsession-with-mayonnaise>.
- Sulis, Emilio, Delia Farías, Paolo Rosso, Viviana Patti, and Giancarlo Ruffo. 2016. Figurative Messages and Affect in Twitter: Differences between #irony, #sarcasm and #not. *Knowledge-Based Systems* 108: 132–143.

- Taboada, Maite, Julian Brooke, Milan Tofiloski, Kimberly Voll, and Manfred Stede. 2011. Lexicon-based Methods for Sentiment Analysis. *Computational Linguistics* 37 (2): 267–307.
- Tutubalina, Elena. 2015. Target-based Topic Model for Problem Phrase Extraction. In *European Conference on Information Retrieval*, 271–277. Cham: Springer.
- Vepsäläinen, Tapio, Hongxiu Li, and Reima Suomi. 2017. Facebook likes and Public Opinion: Predicting the 2015 Finnish Parliamentary Elections. *Government Information Quarterly* 34 (3): 524–532.
- Vilares, David, Thelwall Mike, and Miguel Alonso. 2015. The Megaphone of the People? Spanish SentiStrength for Real-Time Analysis of Political Tweets. *Journal of Information Science* 41 (6): 799–813.
- Volkova, Svitlana, and Eric Bell. 2016. Account Deletion Prediction on RuNet: A Case Study of Suspicious Twitter Accounts Active During the Russian-Ukrainian Crisis. *Proceedings of NAACL-HLT*, 1–6.
- Volkova, Svitlana, Glen Coppersmith, and Benjamin Van Durme. 2014. Inferring User Political Preferences from Streaming Communications. *Proceedings of ACL-2014*, vol. 1, 186–196.
- Volkova, Svitlana, Yoram Bachrach, Michael Armstrong, and Vijay Sharma. 2015. Inferring Latent User Properties from Texts Published in Social Media. *Proceedings of AAAI-2015*, 4296–4297.
- Whalley, Zita. 2018. Why Russians are Obsessed with Mayonnaise? <https://theculture-trip.com/europe/russia/articles/why-russians-are-obsessed-with-mayonnaise/>.
- Wiegand, Michael, Alexandra Balahur, Benjamin Roth, Dietrich Klakow, and Andres Montoyo. 2010. A Survey on the Role of Negation in Sentiment Analysis. *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, 60–68. Association for Computational Linguistics.
- Wilson, Theresa, and Dan Sperber. 2007. On Verbal Irony. *Irony in Language and Thought*: 35–56.
- Wilson, Theresa, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, 347–354.
- Zefirova, Tatyana, and Natalia Loukachevitch. 2019. Irony and Sarcasm Expression in Twitter. *EPiC Series in Language and Linguistics. Proceedings of Third Workshop “Computational Linguistics and Language Science”*, vol. 4, 45–49.

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

