

ИЗВЛЕЧЕНИЕ ИНФОРМАЦИИ ИЗ КЛИНИЧЕСКИХ ТЕКСТОВ НА РУССКОМ ЯЗЫКЕ

Шелманов А. О. (shelmanov@isa.ru),

Смирнов И. В. (ivs@isa.ru)

Институт системного анализа Российской
академии наук, Москва, Россия

Вишнёва Е. А. (vishneva@nczd.ru)

Научный центр здоровья детей, Москва, Россия

Ключевые слова: обработка клинических текстов, извлечение информации, клинические тексты, извлечение атрибутов заболеваний, медицинские тексты

INFORMATION EXTRACTION FROM CLINICAL TEXTS IN RUSSIAN

Shelmanov A. O. (shelmanov@isa.ru),

Smirnov I. V. (ivs@isa.ru)

Institute for Systems Analysis of Russian Academy of Sciences,
Moscow, Russia

Vishneva E. A. (vishneva@nczd.ru)

Scientific Centre of Children Health, Moscow, Russia

We present and evaluate the pipeline for processing of clinical notes in Russian. The paper addresses the tasks of drug identification and disease template filling, which are related to entity recognition and relation extraction. The disease template filling consists in recognition of disease mentions in text, mapping them to concepts of a thesaurus, and discovering their attributes. Discovering attributes means identifying corresponding spans in text, linking them to diseases, and normalizing them i.e. determining their generalized meaning from a predefined set. We implemented tools for determining the following attributes of disease mentions: negation; the flag indicating the disease mention is not related to a patient; severity; course; and body site. For different tasks, we used different techniques: rule-based patterns and several supervised machine-learning methods. Since there were no annotated corpora of clinical notes in the Russian language available for research purposes, we annotated a dataset, which we used for training and evaluation of the developed tools. The created corpus is available for researchers through the data use agreement.

Keywords: clinical text processing, information extraction, annotated corpus, clinical narrative, disease template filling, medical text, EHR

1. Introduction

A vast amount of clinical data is stored as free text. Electronic health records of medical facilities accumulate radiology, echocardiography, and electrocardiogram reports, anamnesis, results of ultrasound diagnostics, discharge summaries, and many other types of notes related to patient healthcare history that are written in a natural language. This is a very rich knowledge source, which is still difficult to exploit because of its unstructured nature. Reworking it into computable form can benefit the biomedical research, patient medical history management, and eventually improve the healthcare. Although there are many natural language processing techniques developed for information and knowledge extraction, the specificity of clinical narrative and tasks arising in the medical domain facilitate the development of specialized methods, language resources, and tools. It is a promising and fruitful scientific direction in natural language processing.

Much of the research in this direction is focused on processing English clinical texts. However, it is also important to create tools and resources for other languages. In the current research, we present and evaluate the pipeline for processing of clinical notes in Russian. The paper addresses the tasks of drug identification and disease template filling, which are related to entity recognition and relation extraction. The disease template filling consists in recognition of disease mentions in text, mapping them to concepts of a thesaurus, and discovering their attributes. Discovering attributes means identifying corresponding spans in text, linking them to diseases, and normalizing them i.e. determining their generalized meaning from a predefined set. We implemented tools for determining the following attributes of diseases: negation; the flag indicating the disease mention is not related to a patient; severity; course; and body site. For different tasks, we used different techniques: rule-based patterns and several supervised machine-learning methods. Since there were no annotated corpora of clinical notes in the Russian language available for research purposes, we annotated a dataset, which we used for the training and evaluation of the developed tools.

2. Related work

Natural language processing of medical texts is a rapidly developing research area. A big number of challenges conducted in the last few years that are devoted to the problems of clinical and biomedical information retrieval and text processing reflects the growing interest in this area for academic community. We briefly review some of them related to information extraction from English clinical texts. 2012 i2b2 /VA Challenge was focused on temporal relation extraction from clinical narratives (Sun et al., 2013). Pilot ShARe/CLEF eHealth 2013 Evaluation Lab set tasks of identifying in clinical reports disease mentions and acronym abbreviations as well as mapping them to thesaurus (Suominen et al., 2013). On ShARe/CLEF eHealth 2014 Evaluation Lab these tasks were extended to disease template filling—identification of disease mentions and their attributes (severity, course, body site etc.). SemEval 2014 Task 7 was similar to CLEF eHealth 2013 and was aimed at disease and acronym/abbreviation identification and

normalization¹. Sem Eval 2015 Task 6 was about temporal information extraction from clinical texts². SemEval-2015 Task 14 was similar to CLEF eHealth 2014 disease template filling task³. We should also note CLEF eHealth 2015 challenge that although is not directly related to clinical text processing, but shows that efforts are not limited to medical information extraction for English language. It introduces the task of named entity recognition in French biomedical texts⁴. This initiative became possible because of recent creation of French annotated biomedical corpus Quaero (Névéol et al., 2014).

The problems of processing clinical notes in Russian we focus on in the current work are largely similar to the tasks of entity recognition and disease template filling introduced by the CLEF eHealth 2014 challenge. The majority of the CLEF eHealth participants applied to this task both the rule-based approaches and supervised machine-learning techniques for extracting different attributes, e.g. (Hamon et al., 2014; Liu and Ku, 2014; Huynh and Ho, 2014). (Ramanan and Nathan, 2014) applied to this task purely rule-based methods. In (Mkrtychyan and Sonntag, 2014), approach based on deep dependency parsing was implemented. Several participants adapted and expanded the cTAKES (Savova et al., 2010) medical text-processing framework (Sequeira et al., 2014; Johri et al., 2014). The participants commonly exploited the MetaMap (Aronson and Lang, 2010) and NegEx (Chapman et al., 2013) tools for mapping found terms to thesaurus concepts and negation detection correspondingly.

Besides, the problem of disease template filling was recently addressed in (Dligach et al., 2014). The severity and body site attributes were linked to disease annotations with SVM classifier. The researchers tested several kernels including the novel tree-based kernel, and found that the common radial-based kernel was suitable for the aforementioned tasks. The evaluation against rule-based baselines showed the advantage of machine learning techniques.

Since there was no basic toolchain for processing clinical texts in Russian, we had to create it from the scratch. Our medical term (diseases, drugs, symptoms, body sites) identification and normalization module is an analogue of MetaMap. The module for disease negation detection and the module for detection that a disease is not related to a patient apply approaches implemented in NegEx. Although our modules are based on somewhat similar approaches to those implemented in the tools for English text processing, we made some modifications in them that take into account peculiarities of the Russian language. In modules for discovering severity and course attributes of disease mentions as well as for linking body sites to disease mentions, we implemented the state-of-the-art machine learning techniques. The annotated corpus we are currently developing contains Russian clinical free-text notes annotated with treatment, symptom, drug, body sites, disorder / disease mentions, their attributes and relations. The closest analogous for English are the corpus of the

¹ <http://alt.qcri.org/semeval2014/task7/>

² <http://alt.qcri.org/semeval2015/task6/>

³ <http://alt.qcri.org/semeval2015/task14/>

⁴ <https://sites.google.com/site/clefehealth2015/>

Shared Annotated Resources (ShARe) initiative (Pradhan et al., 2015) and the corpus of Strategic Health IT Advanced Research Project: Area 4 (SHARPn)⁵.

3. Methods for processing of clinical texts in Russian

The developed clinical text-processing pipeline begins with basic NLP analysis. For tokenization, sentence splitting, part-of-speech tagging, and lemmatization we used well-known pipeline for Russian from AOT.ru (Sokirko, 2001). The dependency syntactic parsing was performed by MaltParser (Nivre et al., 2007) trained on SynTagRus (Apresjan et al., 2005) with configuration described in (Nivre et al., 2008; Sharoff and Nivre, 2011; Smirnov et al., 2014). We find the performance of both of these tools on clinical texts satisfactory. The next step in the pipeline is medical term identification and normalization. It is followed by the disease/symptom negation detection and determining whether disease mentions are related to a patient or not. We will refer to the latter task as “not patient” flag detection. The pipeline concludes by discovering severity and course attributes of disease mentions and linking body sites to disease mentions.

3.1. Identifying medical terms in text and mapping them to thesauri

To solve many tasks of clinical and biomedical text processing it is necessary to perform the basic identification of medical terms in text, term normalization, and mapping them to semantic types. For these tasks, we developed a heuristic- and thesaurus-based module. We used two thesauri: UMLS Metathesaurus (Schuyler et al., 1993) for disease, symptom, and body site identification; a thesaurus based on State Register of Drugs (SRD)⁶ for the drug identification.

UMLS Metathesaurus is an extensive compendium of medical lexicons, classifications, code sets, and thesauri. The main feature of it is that the concepts with the similar sense from different knowledge sources are mapped to the single CUI (concept unique identifier) in the metathesaurus. UMLS also provides Semantic Network that maps concepts into one or more coarse-grained semantic types (McCray, 1989). The only Russian thesaurus present in UMLS is MeSHRUS⁷.

SRD is a database of all drugs officially registered and allowed for sale in Russia. Records in SRD among other information contain trade names of drugs and their active chemicals. Since many drugs have different trade names but almost the same compound and can substitute each other in medication, we preprocessed the database and grouped drugs with similar active chemicals into concepts of thesaurus.

The method implemented in our module was mainly inspired by MetaMap—a widely used linguistically motivated tool for mapping terms from medical texts to concepts

⁵ http://informatics.mayo.edu/sharp/index.php/Main_Page

⁶ <http://grls.rosminzdrav.ru/>

⁷ <http://www.nlm.nih.gov/research/umls/sourcereleasedocs/current/MSHRUS/>

in UMLS Metathesaurus. This tool can be applied only to English text because it strictly relies on handcrafted heuristics and English lexis. The developed module for processing texts in Russian is also rule-based. It generates extensive amount of term variants from text expressions, performs fuzzy comparison between the variants and the thesaurus terms, then ranks the variants by heuristically reasoned score, and picks the most confident ones.

In the first step, the parser scans the text for single keywords from the thesaurus using its reverted index, which is preliminarily built by mapping lemmatized tokens of thesaurus terms to thesaurus concepts. Stop-words like particles, punctuation, prepositions are pruned from the index. The found keywords become the seed term variants for the second step—generation of complex variants.

The generation procedure constructs new term variants from a keyword by expanding it with nodes of the syntactic tree and with words from linear context around the keyword. The procedure is driven by heuristics and constrained by parser parameters (maximum syntax tree depth, maximum number of nodes above the keyword, window of the linear context etc.).

The third step is the assessment of how the generated variants correspond to the thesaurus terms. For each token in the variant, the parser determines the sets of thesaurus terms that correspond to the token and unite them in a single set. Then parser assesses the similarity between the variant and terms from the set. The assessment is a value between 0 (the weakest match) and 1 (the strongest match), which is computed as a linear convolution of three components: “lexical involvement”, “centrality”, and “cohesiveness” (values between 0 and 1). Although the sense of these components is somewhat similar to analogues implemented in the MetaMap, we compute them in the different way.

The “lexical involvement” assesses how the tokens of the variant correspond to the tokens of the thesaurus term. The component ignores order of tokens and considers terms as a bag of words. The lexical involvement is a weighted harmonic mean F_β of two values: the total weight of tokens of a variant among tokens of a term from the thesaurus and the total weight of tokens of a term from the thesaurus among tokens of a variant. Unlike the MetaMap, which convolves these values linearly, we used harmonic mean to combine them because it more strictly penalizes the component if any of them has a small value.

The “centrality” shows whether the most significant token of a variant from syntax perspective (the head of the phrase) is present in the thesaurus term. It is 1 if the syntactic head of the biggest phrase in the variant is present in the term of the thesaurus and 0 otherwise. This component can prune variants with many less significant token matches that lack the most important match of the syntactic head.

The “cohesiveness” extends the idea of “lexical involvement” to syntactically connected phrases. It is a weighted harmonic mean of two values “coherence” and “variant coverage”. The “coherence” is a maximum ratio of a number of syntactically connected tokens in the variant that participate in the match with the thesaurus term and a number of tokens in the variant. This component can prune variants with matches that are not syntactically linked and therefore are not related.

Finally, variant—thesaurus term pairs are selected by a threshold and filtered by heuristics. The chosen variants are mapped to the semantic types according to their identifier with the biggest score.

3.2. Negation and “not patient” flag detection

It is crucial to know whether the found disease or symptom mentions are negated in text and whether they are associated with the patient or with another person, e.g. patient’s relative, since many clinical notes also contain information about patient’s heredity.

For negation detection, medical text processing systems usually apply simple approach based on pattern matching. For example, NegEx implements an algorithm that searches for a list of patterns in a linear context in a window around a disease mention. Despite the simplicity of the algorithm, it proved itself robust with moderate performance and was adopted in cTAKES.

Besides the aforementioned strategy, we also implemented pattern search in a syntax tree, which showed somewhat better performance. The latter approach was integrated into the final pipeline. The developed negation detection module searches for the following patterns: particle “не” (“not”) syntactically depends on one of the tokens of the disease/symptom; particle “не” (“does not”) syntactically depends on predicate (single predicate word of complex predicate with auxiliary verb) that governs a token of the disease/symptom term; particle “нет” (“no”) governs a token from the disease/symptom term; a token of the disease/symptom term is governed by negation predicate, e.g. “отсутствует” (“is absent”); particle “нет” (“no”) immediately follows a disease/symptom mention.

Module that detects “not patient” flag mainly searches for mentions of patient relatives and patterns that associate them with disease terms in text: “y” + “relative mention” syntactically connected to or precedes disease term in a sentence; “наследственность” (“heredity”) precedes disease mention in a sentence.

3.3. Discovering severity, course, and body site attributes of diseases

Severity attributes capture and normalize text cues that clarify whether the disease mentioned in text is “slight”, “moderate”, or “severe”. Course attributes capture and normalize text cues about disease progress: “worsened”, “changed”, “improved”, “resolved”. Body site attributes determine which body part disease mentions are associated with in a clinical text. In example: “*Диагностирована астма, atopическая, легкое персистирующее течение, ремиссия. ... Имеется ангиопатия сетчатки.*” (“*Diagnosed asthma, atopic, mild persistent, during remission. ... Retinal angiopathy is present.*”). The string “*легкое персистирующее течение*” (“*mild persistent*”) should be captured in the severity attribute of the disease mention “*астма, atopическая*” (“*asthma, atopic*”) and normalized as “slight”; the string “*ремиссия*” (“*during remission*”) should be captured in the course attribute of the same disease mention and normalized as “improved”; the string “*сетчатки*” (“*retinal*”) should be marked as body site and linked to the disease mention “*ангиопатия*” (“*angiopathy*”). The common practical usage of the extracted information is summarizing electronic health records for information or analytical systems in terms of standardized format like Clinical Document

Architecture⁸. For the tasks of discovering severity, course, and body site attributes, we implemented several modules based on supervised machine learning methods.

The modules that discover severity and course attributes are almost the same. They consist of two separate submodules for the attribute span identification with linking it to the corresponding disease mention and for the attribute normalization. For the given disease mention, the submodule for the attribute span identification scans tokens of a sentence, in which the corresponding disease mention is located, and applies to them binary classifier that predicts whether token is an attribute related to the disease mention or not. When every token is classified, the tokens marked by the same attribute are grouped into continuous annotations. We used the following lexical and syntactic features: lemmas and postags of tokens in a window around the classified token; whether the classified token syntactically depends from the disease term; distance between the classified token and the disease term; relative position of the token regarding to the given disease mention; number of disease annotations between the disease mention and the token. When attribute spans are identified, another submodule with a separate classifier normalizes them. The feature set of the latter classifier is composed of token lemmas lying in the corresponding span of severity or course annotation represented as a bag of words.

In the task of discovering body site attributes, the spans that represent body parts are identified by thesaurus-based parser described in section 3.1. Therefore, only linking these spans to disease mentions is required. The module for linking body sites to diseases scans all body site—disease pairs within a sentence and analyzes them with a binary classifier. The features of the classifier include distance in tokens between a disease mention and a body site, whether they are syntactically linked, whether they are attached to the same word (e.g., predicate), the postag of this word, the number of disease mentions between the given disease mention and the body site.

We tested several classifiers for each of the tasks and subtasks: linear SVM, SVM with radial basis kernel, random forest, and AdaBoost. The parameter tuning was performed on the developed corpus via cross-validation.

4. Annotated corpus of clinical notes in Russian

In conjunction with specialists of Scientific Center of Children Health (SCCH)⁹ we created annotated corpus of clinical free-text notes in Russian. The corpus is based on medical histories of more than 60 SCCH patients with allergic and pulmonary disorders and diseases. It comprises discharge summaries, radiology, echocardiography, ultrasound diagnostics reports, recommendations, and other records created by different physicians. The documents in the corpus were de-identified: all names were removed and dates were altered. With the help of SCCH experts, we developed an annotation scheme and a guideline. The scheme encompasses span annotations: “Disease”, “Symptom”,

⁸ http://www.hl7.org/implement/standards/product_brief.cfm?product_id=7

⁹ <http://www.nczd.ru/eng/>

“Drug”, “Treatment”, “Body location”, “Severity”, “Course”; attributes: “Negation” (for “Disease”, “Symptom”), “NotPatient” (for “Disease”, “Symptom”, “Treatment”), “Degree” (for “Severity”), “CType” (for “Course”); relations between “Severity” and “Disease”, “Course” and “Disease”, “Body location” and “Disease”; and other annotations.

The corpus was annotated and verified by physicians. For annotation purposes, we used Brat—the web-based tool, which was originally created for BioNLP challenge¹⁰. The Fig. 1 illustrates a fragment of an annotated text.

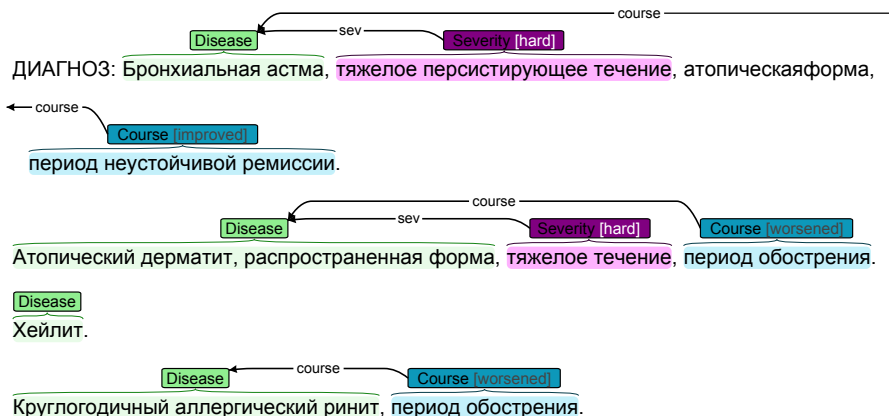


Fig. 1. Example of annotated text

The corpus currently contains more than 112 fully annotated texts with almost 45,000 tokens. There are more than 7,600 annotated entities and more than 4,000 annotated attributes and relations. The work on the corpus is still in progress: we are adding new texts and planning to expand the annotation scheme.

The corpus is freely available for the research community through the data use agreement¹¹. However, the human subjects training certificate¹² would be required for access to it since the corpus contains the medical data of real patients.

5. Experiments

5.1. Disease and drug identification performance

The testing of medical term identification module was performed on a randomly selected holdout consisting of 30 texts; the rest of the corpus was used for parameter tuning. We calculated relaxed versions of precision, recall, and F_1 score. In the relaxed assessment, a span overlapping a gold standard span is considered correct.

¹⁰ <http://brat.nlplab.org/>

¹¹ <http://nlp.isa.ru/datasets/clinical>

¹² <https://phrp.nihtraining.com/users/login.php>

Especially for the disease identification task, we prepared two baselines. The first baseline marks in text all words of thesaurus concepts related to “disease” semantic type. This baseline tends to maximum recall. The second baseline marks in text only token chains that exactly match a whole bag of words of a thesaurus concept. This baseline tends to maximum precision. Table 2 presents performance of the developed module and the baselines.

Table 2. Performance of the disease identification module and the baselines

Module	Recall,%	Precision,%	F ₁ -score,%
Disease identification	72.8	95.1	82.4
Baseline 1	84.9	9.3	16.7
Baseline 2	69.8	99.2	81.9

As expected, the implemented module has better overall performance than the baselines, however, it appeared to be very close to the performance of the baseline 2. The developed method should more significantly outmatch the baseline 2 in the task of mapping terms to concepts because the baseline 2 often maps terms to the very general concepts although they should be mapped to the more specific ones. For example, the module finds term “аллергический ринит” (“allergic rhinitis”) in text. This string cannot be mapped exactly to the bag of words of any thesaurus term, so the baseline 2 will not find it. However, it will find “ринит” (“rhinitis”) and map it to the more general concept. Both answers will be considered as a match for the disease identification task, however, in the concept-mapping task, the answer of the module will be the only correct one.

For the evaluation of the drug identification performance, we used exactly the same framework as for the disease identification. The precision was 84.3, the recall was 74.6, and the F₁-score was 79.2. The recall is somewhat lower than expected because annotators marked not only registered drugs but also mentions of therapeutic cosmetics (e.g., anti-allergic cream). We also note that some false negatives were due to corpus texts contained contractions and general names of drugs like “пенициллин” (“penicillin”), which is not present in SRD, instead of full name like “бензилпенициллин” (“benzathine penicillin”), which is in SRD.

5.2. Symptom negation and “not patient” flag detection performance

We found that the annotated corpus contains just a few negations of disease mentions, much more negations are related to symptoms. Therefore, instead of disease negation detection we only evaluated the performance of the symptom negation detection.

The number of negation and “not patient” flag annotations in corpus is relatively small—both around 100. To make the test set representative, we had to evaluate negation and “not patient” flag detection using the whole corpus rather than the selected holdout. Since the developed modules are rule-based, such approach has a minor effect on the performance assessment bias. In the evaluation, we only took into account attributes of symptom and disease mentions identified by our system to exclude false negatives of the term identification module (table 3).

Table 3. Performance of the symptom negation and “not subject” detection

Module	Recall,%	Precision,%	F ₁ -score,%
Symptom negation	98.7	95.3	97.0
Disease “not patient”	90.9	96.8	93.8

The obtained results show that a rather simple approach and few patterns can cover the most cases of negation and “not patient” attributes in clinical texts. However, we admit that the current test set is small and not very representative for the perfect evaluation.

5.3. Performance of discovering severity, course, and body site attributes of diseases

For all evaluations, we used 5-fold cross validation on the annotated corpus.

We separately evaluated performance of the identification and normalization of severity and course attributes. To exclude the false negatives of the disease identification module, we only took into account course and severity annotations that are related to disease mentions identified by our system. For evaluation of severity and course identification, similarly to the evaluation of medical term identification, we calculated relaxed versions of precision, recall, and F₁-score (tables 4 and 5).

Attribute normalization is a multilabel classification task without “empty” class. Therefore, evaluation of severity and course normalization was performed via accuracy (table 6).

The performance of linking body sites to diseases was evaluated via relaxed precision, recall, and F₁-score. Only disease mentions that were identified by the system were taken into account (table 7).

Table 4. Performance of severity identification

Classifier	Recall,%	Precision,%	F ₁ -score,%
Linear SVM	99.2	41.7	58.6
RBF SVM	95.0	80.8	87.1
Random forest	93.6	82.6	87.5
AdaBoost (Dec. tree)	97.3	75.2	84.7

Table 5. Performance of course identification

Classifier	Recall,%	Precision,%	F ₁ -score,%
Linear SVM	92.3	99.2	95.7
RBF SVM	88.3	99.3	93.4
Random forest	88.3	99.3	93.4
AdaBoost (Dec. tree)	90.0	98.4	93.9

Table 6. Performance of severity and course normalization

Module	Classifier	Accuracy,%
Severity normalization	Linear SVM	88.4
	RBF SVM	88.0
	Random forest	89.3
	AdaBoost (Dec. tree)	89.8
Course normalization	Linear SVM	89.4
	RBF SVM	91.4
	Random forest	92.7
	AdaBoost (Dec. tree)	91.4

Table 7. Performance of linking body sites to diseases

Classifier	Precision, %	Recall, %	F ₁ -score, %
Linear SVM	85.4	77.5	81.0
RBF SVM	91.4	76.6	83.3
Random forest	86.6	75.8	80.8
AdaBoost (Dec. tree)	84.0	76.6	79.9

Experiments with different classifiers showed statistical difference between their results only in the task of severity identification. This can be because of importance of word collocations that are differently modeled by these classifiers. The developed modules represent a solid baseline. However, there is still a big space for improvement, which can be achieved by developing richer set of features and applying new machine learning methods. We consider that the obtained results can be a useful landmark for the future research.

6. Conclusion and future work

We presented the pipeline for processing of clinical texts in Russian and the corpus of annotated clinical notes. We evaluated the pipeline and showed that it can successfully solve the key tasks of information extraction from clinical narrative. In the ongoing work, we are expanding annotated corpus to make it more representative and suitable for training machine-learning algorithms, as well as for reliable testing of clinical text processing methods and tools. In the future work, we are planning to upgrade corpus annotation scheme and include in the pipeline modules for discovering treatments with modifiers and temporal expressions related to disease mentions. We also are planning to apply the developed pipeline for the high-level task of clinical information retrieval and clinical data analysis.

Acknowledgments

We are grateful to experts of Scientific Center of Children Health for help on annotating corpus of clinical texts in Russian. This work was supported by RFBR, project 13-04-12062.

References

1. *Apresjan J. D., Boguslavskij I. M., Iomdin B. L., Iomdin L. L., Sannikov A. V., Sannikov V. G. and Sizov L. L.* (2005), Syntactically and semantically annotated corpus of Russian language: Present state and perspectives [Sintaksicheski i semanticheski annotirovannyj korpus russkogo jazyka: sovremennoe sostojanie i perspektivy], National Corpus of Russian Language: 2003–2005 [Natsional'nyj korpus russkogo jazyka: 2003–2005], pp. 193–214, (in Russian)
2. *Aronson, A. R. and Lang, F.-M.* (2010), An overview of MetaMap: historical perspective and recent advances, *Journal of the American Medical Informatics Association (JAMIA)*, (3), Vol. 17, pp. 229–236
3. *Bos, L. and Donnelly, K.* (2006), SNOMED-CT: The advanced terminology and coding system for eHealth, *Studies in health technology and informatics*, Vol. 121, pp. 279–290
4. *Chapman, W. W., Hilert, D., Velupillai, S., Kvist, M., Skeppstedt, M., Chapman, B. E., Conway, M., Tharp, M., Mowery, D. L. and Deleger, L.* (2013), Extending the NegEx lexicon for multiple languages, *Studies in health technology and informatics*, Vol. 192, pp. 677–681
5. *Dligach, D., Bethard, S., Becker, L., Miller, T. A. and Savova, G. K.* (2014), Discovering body site and severity modifiers in clinical texts, *Journal of the American Medical Informatics Association (JAMIA)*, pp. 448–454
6. *Hamon, T., Grouin, C. and Zweigenbaum, P.* (2014), Disease and Disorder Template Filling using Rule-based and Statistical Approaches, In *CLEF (Working Notes)*, pp. 79–90
7. *Huynh, H. N. and Ho, S. L. V. B. Q.* (2014), ShARE/CLEFeHealth: A Hybrid Approach for Task 2, In *CLEF (Working Notes)*, pp. 103–110
8. *Johri, N., Niwa, Y. and Chikka, V. R.* (2014), Optimizing Apache cTAKES for Disease/Disorder Template Filling: Team HITACHI in the ShARE/CLEF 2014 eHealth Evaluation Lab, In *CLEF (Working Notes)*, pp. 111–123
9. *Lipscomb, C. E.* (2000), Medical subject headings (MeSH), *Bulletin of the Medical Library Association*, Vol. 88, pp. 265–266
10. *Liu, Y.-C. and Ku, L.-W.* (2014), CLEFeHealth 2014 Normalization of Information Extraction Challenge using Multi-model Method, In *CLEF (Working Notes)*, pp. 124–132
11. *McCray, A. T.* (1989), The UMLS Semantic Network. In *Proceedings of the Annual Symposium on Computer Application in Medical Care*, pp. 503–507, American Medical Informatics Association
12. *Mkrtchyan, T. and Sonntag, D.* (2014), Deep Parsing at the CLEF2014 IE Task, In *CLEF (Working Notes)*, pp. 138–146

13. *Névéol, A., Grouin, C., Leixa, J., Rosset, S. and Zweigenbaum, P.* (2014), The Quaero French medical corpus: A resource for medical entity recognition and normalization, In Proceedings of LREC BioTxtM 2014 Fourth Workshop on Building and Evaluating Resources for Health and Biomedical Text Processing
14. *Nivre, J., Boguslavsky, I. M. and Iomdin, L. L.* (2008), Parsing the SynTagRus treebank of Russian, In Proceedings of the 22nd International Conference on Computational Linguistics, pp. 641–648, Association for Computational Linguistics
15. *Nivre, J., Hall, J., Nilsson, J., Chanev, A., Eryigit, G., Kübler, S., Marinov, S. and Marsi, E.* (2007), MaltParser: A language-independent system for data-driven dependency parsing, *Natural Language Engineering*, (2), Vol. 13, pp. 95–135
16. *Pradhan, S., Elhadad, N., South, B. R., Martinez, D., Christensen, L., Vogel, A., Suominen, H., Chapman, W. W. and Savova, G.* (2015), Evaluating the state of the art in disorder recognition and normalization of the clinical narrative, *Journal of the American Medical Informatics Association (JAMIA)*, (1), Vol. 22, pp. 143–154
17. *Ramanan, S. V. and Nathan, P. S.* (2014), Cocoa: Extending a Rule-based System to Tag Disease Attributes in Clinical Records, In CLEF (Working Notes), pp. 150–155
18. *Savova, G. K., Masanz, J. J., Ogren, P. V., Zheng, J., Sohn, S., Kipper-Schuler, K. C. and Chute, C. G.* (2010), Mayo clinical Text Analysis and Knowledge Extraction System (CTAKES): Architecture, component evaluation and applications, *Journal of the American Medical Informatics Association (JAMIA)*, (5), Vol. 17, pp. 507–513
19. *Schuyler, P. L., Hole, W. T., Tuttle, M. S. and Sherertz, D. D.* (1993), The UMLS Metathesaurus: Representing different views of biomedical concepts, *Bulletin of the Medical Library Association*, (2), Vol. 81, 217–222
20. *Sequeira, J., Miranda, N., Goncalves, T. and Quaresma, P.* (2014), TeamUEvora at CLEF eHealth 2014 Task2a, In CLEF (Working Notes), pp. 156–166
21. *Sharoff S. and Nivre J.* (2011), The proper place of men and machines in language technology: Processing Russian without any linguistic knowledge, *Computational Linguistics and Intellectual Technologies: Proceedings of the International Conference “Dialog 2011”*, pp. 591–604
22. *Smirnov, I. V., Shelmanov, A. O., Kuznetsova, E. S. and Hramoin, I. V.* (2014), Semantic-syntactic analysis of natural languages. Part II. Method for semantic-syntactic analysis of texts [Semantiko-sintaksicheskij analiz estestvennyh jazykov Chast’ II. Metod semantiko-sintaksicheskogo analiza tekstov], *Artificial Intelligence and Decision Making [Iskusstvennyj intellekt i prinjatje reshenij]*, (1), pp. 95–108, (in Russian)
23. *Sokirko, A.* (2001), A short description of Dialing Project, available at: <http://www.aot.ru/>
24. *Sun, W., Rumshisky, A. and Uzuner, O.* (2013), Evaluating temporal relations in clinical text: 2012 i2b2 challenge, *Journal of the American Medical Informatics Association (JAMIA)*, (5), Vol. 20, pp. 806–813
25. *Suominen, H., Salanterä, S., Velupillai, S., Chapman, W. W., Savova, G., Elhadad, N., Pradhan, S., South, B. R., Mowery, D. L., Jones, G. J. et al.* (2013), Overview of the ShARe/CLEF eHealth evaluation lab 2013, In *Information Access Evaluation. Multilinguality, Multimodality, and Visualization*, Lecture Notes in Computer Science, Vol. 8138, pp. 212–231, Springer