

Original Research

Multimodal model with text and drug embeddings for adverse drug reaction classification

 Andrey Sakhovskiy^{a,b}, Elena Tutubalina^{a,c,d,*}
^a Kazan Federal University, 18 Kremlyovskaya street, Kazan, 420008, Russian Federation

^b Lomonosov Moscow State University, 1 Leninskie gory, Moscow, 119991, Russian Federation

^c Sber AI, 19 Vavilova St., Moscow, 117997, Russian Federation

^d National Research University Higher School of Economics, 11 Pokrovsky Bulvar, Moscow, 109028, Russian Federation


ARTICLE INFO

Keywords:

Natural language processing

Social media

Adverse drug reactions

ABSTRACT

In this paper, we focus on the classification of tweets as sources of potential signals for adverse drug effects (ADEs) or drug reactions (ADRs). Following the intuition that text and drug structure representations are complementary, we introduce a multimodal model with two components. These components are state-of-the-art BERT-based models for language understanding and molecular property prediction. Experiments were carried out on multilingual benchmarks of the Social Media Mining for Health Research and Applications (#SMM4H) initiative. Our models obtained state-of-the-art results of 0.61 F_1 -measure and 0.57 F_1 -measure on #SMM4H 2021 Shared Tasks 1a and 2 in English and Russian, respectively. On the classification of French tweets from SMM4H 2020 Task 1, our approach pushes the state of the art by an absolute gain of 8% F_1 . Our experiments show that the molecular information obtained from neural networks is more beneficial for ADE classification than traditional molecular descriptors. The source code for our models is freely available at https://github.com/Andoree/smm4h_2021_classification.

1. Introduction

The popularity of social media as a source for health-related information has increased tremendously in the past decade. One of the well-studied research areas is pharmacovigilance from social media data that focuses on discovering adverse drug effects (ADEs) from user-generated texts. ADEs¹ are unwanted negative effects of a drug, in other words, harmful and undesired reactions due to its intake.

In recent years, researchers have increasingly applied neural networks, including Bidirectional Encoder Representations from Transformers (BERT) [1], to ADE detection from texts [2–5]. This is directly related to the creation of annotated multilingual corpora and Social Media Mining for Health (#SMM4H) shared tasks [2,3]. Given a tweet, participants of this shared task are required to detect whether the tweet contains a mention of ADE using natural language processing (NLP) techniques. However, these studies mostly share the same limitations: models consider textual information only without leveraging drug structure. In particular, [6,7] utilized transformer-based classifiers with ensemble modeling and undersampling ranking first on SMM4H 2020 task 1a and 1b, respectively. On the other hand, studies from cheminformatics and drug discovery areas have focused on the

prediction of the side effects of a given drug [8–11]. These studies utilized supervised models trained and evaluated on a database of marketed drugs and ADRs the Side Effect Resource (SIDER) [12] from the MoleculeNet benchmark [13]. Each chemical structure is encoded with molecular descriptors or neural representations that are usually fed to a multi-label classification model. In addition, a number of BERT architectures has been proposed for molecular property prediction such as SMILES-BERT [11], MolBERT [14], ChemBERTa [15] with fine-tuning on the SIDER dataset.

Inspired by multimodal studies, we propose a novel method to utilize both textual and molecular information for ADE classification. We study the impact of using different molecular representation approaches, including traditional molecular descriptors calculated with Mordred [16] and BERT-based encoders ChemBERTa and MolBERT. We explore two strategies to fuse drug representations and tweet representations: straightforward concatenation of representations and use of a co-attention mechanism to integrate features of different modalities. Along with the textual information, the incorporation of molecular structure can aid in understanding the relationship between different pharmacological and chemical properties and the occurrence of ADEs.

* Corresponding author at: Kazan Federal University, 18 Kremlyovskaya street, Kazan, 420008, Russian Federation.

E-mail addresses: andrey.sakhovskiy@gmail.com (A. Sakhovskiy), ElVTutubalina@kpfu.ru (E. Tutubalina).

¹ The terms adverse drug effects (ADEs) and adverse drug reactions (ADRs) are often used interchangeably.

<https://doi.org/10.1016/j.jbi.2022.104182>

Received 12 January 2022; Received in revised form 24 May 2022; Accepted 20 August 2022

Available online 30 September 2022

1532-0464/© 2022 Elsevier Inc. All rights reserved.

A preliminary version of this work has appeared in [17]. Compared to the conference version, we have: (1) significantly extended the experimental part of this work to assess the performance of the proposed multimodal model; in particular, we extended experiments to three datasets in English, French, and Russian. Our model achieved state-of-the-art results on SMM4H 2021 Tasks 1a & 2 and SMM4H 2020 Task 1b; (2) extended the description of the experimental datasets for ADE classification; (3) investigated model performance on different drug groups, adding new experimental results and conclusions; (4) performed error analysis and discussed the limitations of our model.

In Section 2 we present related work, Section 3 describes the datasets, Section 4 introduces our approach, Section 5 describes and discusses experimental results, and Section 7 concludes the paper.

2. Related work

2.1. ADE classification

Although there is a wide range of supervised machine learning methods that have been applied to classify user-generated posts in English, a relatively small number of recent studies have been focused on texts in other languages, i.e., Russian and French. In the SMM4H 2020 shared task [3], the precision of optimal ADE classification systems varies across different languages. For English, precision previously has stayed in the range of 0.45–0.65 reaching a score of 0.64 with the winning system, while precision for Russian and French have stayed in the range of 0.34–0.54 and 0.15–0.33, respectively [2]. Most of these models leverage raw text to classify each text or a token in a text to ADE class. [6] pre-trained a large version of RoBERTa [18] on an unlabeled corpus of 6 million tweets in English where each tweet includes drug mention; they fine-tuned this model without any imbalance techniques on SMM4H data. [19] trained a fully-connected network on pre-computed embeddings obtained using a distilled version of multilingual sentence BERT to classify French tweets. [7] used an ensemble of BERT-based models and logistic regression with undersampling to classify Russian tweets. [4] used an ensemble of ten EnRuDR-BERT models [20] pretrained on 5M health-related user posts. Overall, the percentage of teams using Transformer [21] architectures for ADE classification rose from 80% in SMM4H 2020 to 100% in SMM4H 2021 [3].

Our model improves upon the state-of-the-art models in three very critical ways: (1) incorporation of drug representations, (2) extensive experiments on three languages, (iii) analysis of model performance for various Anatomical Therapeutic Chemical (ATC) main groups.

2.2. Multimodal learning on biomedical tasks

A number of studies proposed multimodal deep learning models to combine information from multiple modalities. These models show promising results compared to uni-modal models on protein–protein and drug–drug interaction identification from scientific texts in English [22,23], image captioning [24], image ads understanding [25]. [22] utilize textual sentence representations and representation of the molecular structures of active substances for classification of drug–drug interactions (DDI). The authors proposed a network with two components. The first component of the architecture is a convolutional neural network, where the original text is converted into embedding representations using the word2vec model, which are then combined with the positional embeddings of the two drugs in the text. The second component is a graph convolutional network (GCN). Experiments show a 2.39% increase in the F-measure on the DDI problem compared to methods that do not use graph representations of molecules as features. Similar to the DDI task, [23] investigated multimodal networks on BioInfer HRPD50 datasets for the protein–protein interaction (PPI) task. The authors used BioBERT pre-trained

on large corpora of medical texts (PubMed, PMC) to obtain text representations. Spatial structure in PDB format or FASTA-sequence of each protein is available according to the corresponding identifiers (PDB ID, ensemble ID). To obtain features from FASTA sequences, the nucleotide sequences are first one-hot encoded and serve as the input for three convolutional layers. To obtain structural features, the coordinates of the atoms from the PDB file are converted into an adjacency matrix, and for each atom, its feature vector is calculated. The obtained information is used as the graph representation of the protein, which is fed to GCN. The authors used a transformer layer with attention combining three modalities; the output of this layer serves as an input to the softmax layer for the final classification. On both datasets, the authors achieve state-of-the-art results, in both cases improving the previous F-measure result by about 7% in absolute terms.

3. Data

The ADE classification task involves distinguishing tweets that report an adverse effect of a medication (annotated as “ADE”) from those that do not (annotated as “non-ADE”). The aim of the Social Media Mining for Health Applications (#SMM4H) shared tasks is to take a community-driven approach to addressing NLP challenges of utilizing social media data for health informatics, including informal, colloquial expressions of clinical concepts, noise, data sparsity, ambiguity, and multilingual posts. In 2020, the fifth iteration of the SMM4H shared tasks included Task 2 on automatic classification of multilingual tweets that report adverse effects of a medication. This dataset includes tweets posted in English, French, and Russian [2].

All SMM4H datasets are manually labeled using the same annotation guidelines. According to these guidelines, an ADE or ADR is an effect of the drug that is not desired. This includes a worsening of the indication, drug withdrawal and its associated symptoms, and a loss of effectiveness of the medication. The general rules for marking a tweet as positive for containing a mention of an ADR are: 1. The user states an ADR that resulted from taking the drug; 2. The ADR mentioned had been personally experienced by the user, or by someone known personally to the user; 3. The ADR mentioned is specific, that is it refers to a sign or symptom that can be coded to a medical term even if the mention itself is in colloquial terms; 4. The ADR occurred over the normal course of treatment and is not an effect from taking too much of the drug; 5. The ADR is not the result of taking two, or more, drugs at the same time (drug–drug interactions).

English-language version of the SMM4H task is a long-running task starting from 2016, while French and Russian subtasks were presented for the first time. Within a year, [3] released English and Russian datasets for SMM4H 2021 Task 1a and Task 2, respectively. We note that all Russian texts from SMM4H 2020 were included in the SMM4H 2021 Task 2 train/dev sets. Therefore, we use the latter one for our experiments. Since the organizers of the SMM4H 2021 Shared Task did not publish the original annotated test set, we use 20% of the original train tweets as a test set for our experiments. The data split was performed randomly, but the positive and negative class proportion is kept the same for train and test sets. To explore the effect of drug group distribution among train and test sets, we created an additional ATC-stratified split of the English SMM4H 2021 dataset. In the ATC-stratified data split, the proportion of tweets mentioning an ATC group is the same for train and test sets for a fixed ATC group.

Table 1 presents various statistics of these subsets. Several interesting observations can be made based on Table 1. First, the size of training sets varies significantly: the training set for the English subtask contains 25,678 tweets, while the first-running French set contains only 1941 tweets. Second, the task is especially challenging due to the imbalanced nature of the dataset where the tweets that mention ADE are outnumbered 10:1 to 50:1 by tweets that do not contain ADEs. Third, the number of unique drugs differs significantly for different languages. The definition of drug products and dietary supplements provided by

Table 1
Summary of statistics of experimental datasets.

	Train	Dev	Test	All
SMM4H 2020 Task 1b, French tweets				
# tweets	1941 (1.6%)	485 (1.6%)	607 (1.6%)	3033 (1.6%)
# tweets with ADEs	31	8	10	49
# unique drug ids	175	91	114	219
avg. len	27.54	26.72	27.45	27.39
SMM4H 2021 Task 2, Russian tweets				
# tweets	10,609	1000	9095	20,707
# tweets with ADEs	980 (9.2%)	92 (9.2%)	778 (8.6%)	1850 (8.9%)
# unique drug ids	62	47	41	65
avg. len	21.48	22.51	18.65	20.29
SMM4H 2021 Task 1a, English tweets				
# tweets	13,908	914	3477	18,299
# tweets with ADEs	988 (7.1%)	65 (7.1%)	247 (7.1%)	1300 (7.1%)
# unique drug ids	581	215	340	638
avg. len	19.81	19.75	19.72	19.79
SMM4H 2021 Task 1a, English tweets, ATC-stratified split				
# tweets	13,176	1456	3667	18,299
# tweets with ADEs	968 (7.3%)	102 (7%)	230 (6.3%)	1300 (7.1%)
# unique drug ids	582	234	349	638
avg. len	19.74	20.03	19.89	19.79

Table 2
Most frequent drugs of SMM4H corpora. The number in parenthesis indicates the number of tweets with the drug mention.

Dataset	Most frequent drugs with DrugBank id
SMM4H 2020 Task 1b, French tweets	xanax, DB00404 (929) ventoline, DB01001 (606) prozac, DB00472 (226) levothyrox, DB00451 (203) morphine, DB00295 (101)
SMM4H 2021 Task 2, Russian tweets	paracetamol, DB00316 (6124) umifenovir/arbidol, DB13609 (2363) prozac, DB00472 (2335) xanax, DB00404 (1718) paroxetine, DB00715 (834)
SMM4H 2021 Task 1a, English tweets	adderall, DB01576 (1545) prozac, DB00472 (1109) vyvanse, DB01255 (900) cymbalta, DB00476 (863) xanax, DB00404 (814)

the FDA [26] were used to collect English tweets. The French tweets have been selected based on a list of around 120 drug names. The Russian tweets were collected as follows. For the Russian subset for SMM4H 2020, around 70 drugs most commonly described in English tweets were manually associated with drugs on the Russian market using the State Register of Medicinal Remedies [27]. Additionally, the authors collected tweets for COVID-related drugs such as *paracetamol* and *umifenovir* in 2021 for the test set. Paracetamol, also known as acetaminophen, is used to treat fever and mild pain. Umifenovir, also known as arbidol, is a popular antiviral medication for the treatment of influenza infection used in Russia.

Table 2 shows most frequent drugs of SMM4H corpora. Top 5 drugs in English tweets are parts of the ATC group *Nervous system* (N). Russian tweets include drugs assigned to two groups: *Antiinfectives for systemic use* (J) and *Nervous system* (N). French tweets include drugs assigned to four groups: *Nervous system* (N), *Respiratory system* (R), *Systemic hormonal preparations, excluding sex hormones and insulins* (H), *Alimentary tract and metabolism* (A). We analyzed the number of tweets associated with drugs from different ATC classes for the Russian subset for SMM4H 2021. Fig. 1 presents these results. As expected, some groups such as *Antiinfectives for systemic use* are represented in the test set more significantly than in the train set.

To determine whether an ATC group of a drug mentioned in a text provides an a priori knowledge on how probable is the text to report

an ADE, we calculated the statistics on the percentage of ADE tweets among tweets of each ATC class. The results are presented in Fig. 2. Several important observations have to be done. First, the percentage of tweets mentioning ADE varies significantly among different drug groups. For example, about 2% of English tweets of the “*Antineoplastic and immunomodulating agents*” class mention an adverse effect, whereas more than 10% tweets of the “*Nervous systems*” class mention an ADE. Thus, we can conclude that given a text with drug mention, we are provided with some a priori knowledge on how likely is the text to report an ADE. Second, the proportion of tweets mentioning an ADE is inconsistent among different languages, even for the same ATC group. For example, for “*Musculoskeletal system*” class about 7% of Russian tweets and less than 2% of English tweets report an ADE. Based on the observation, we can hypothesize that drug-related knowledge transfer across multiple languages may fail due to this inconsistency. Besides, drug-related information is not an exhaustive tweet representation itself, and textual context is required.

For our experiments, we use the French sets from the SMM4H 2020 task and English and Russian sets from the SMM4H 2021 tasks. The data preprocessing pipeline is identical for all datasets and is adopted from [4] and includes the following steps: (i) replacement of all URLs with the word “link”; (ii) masking of user mentions with the “@username” placeholder; (iii) replacement of some emojis with a textual representation (e.g., pill and syringe emojis are replaced by “pill” and “syringe” words).

4. Models

In our work, we propose bi-modal ADE text classification models that combine textual and drug modalities. More formally, consider we have a text T that can be represented as a pair of textual modality t and a drug modality that is a set of k_T drug mentions: $T = (t, D_T)$, where $D_T = (d_1, d_2, \dots, d_{k_T})$. To obtain two unimodal representations of T , we sample a random drug d_i from D_T and encode the t and d_i using two encoders: (i) a textual encoder M_{text} and (ii) a drug encoder M_{drug} :

$$u_{text}^T = M_{text}(t) \quad u_{drug}^T = M_{drug}(d_i)$$

where $u_{text}^T \in \mathbb{R}^{d_t}$ is the embedding of textual modality t and $u_{drug}^T \in \mathbb{R}^{d_d}$ is the embedding of drug d_i ; d_t and d_d are the dimensionalities of the obtained uni-modal vector representations. The multimodal binary text classification task can be formulated as:

$$f_{cl}(f_{mod}(u_{text}^T, u_{drug}^T))$$

where $f_{mod} : \mathbb{R}^{d_t+d_d} \rightarrow \mathbb{R}^{d_{bi}}$ is a modality combination function that provides a bi-modal representation of T , $f_{cl} : \mathbb{R}^{d_{bi}} \rightarrow \mathbb{R}$ is a fully connected classification network with sigmoid output activation. d_{bi} is the dimensionality of the bi-modal vector representation.

We use the final embedding of the classification token (“[CLS]” token) as a textual representation of an input text. In the uni-modal approach, the textual representation u_{text}^T is usually directly fed to a classifier, whereas for bi-modal models, we first combine the textual modality with the drug modality. We evaluate 2 ways to combine text and drug modalities: (i) concatenation of text and drug embedding; (ii) the scaled dot-product attention used in Transformer-based models [21]. For the first way, we concatenate the textual embedding and drug embedding:

$$f_{mod}^{concat}(u_{text}^T, u_{drug}^T) = [u_{text}^T \oplus u_{drug}^T]$$

Attention mechanism allows to learn a multimodal embedding of sample T as a linear combination of two modalities u_{text}^T and u_{drug}^T . For text classification, we use the sum of the multimodal embedding and the textual embedding:

$$f_{mod}^{att}(u_{text}^T, u_{drug}^T) = \alpha \cdot u_{text}^T + (1 - \alpha) \cdot u_{drug}^T$$

where α is the weight of textual modality obtained using attention mechanism.

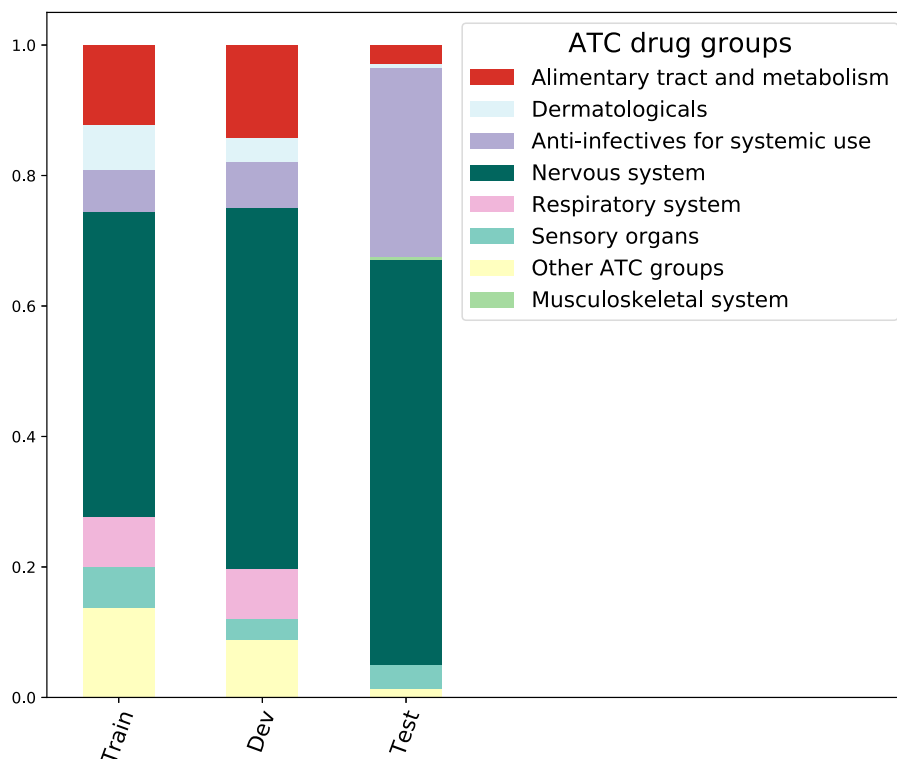


Fig. 1. ATC classes distribution in the Russian corpus for SMM4H 2021.

Recently, a number of BERT-based models were proposed to learn flexible and high-quality molecular representations for drug discovery problems. These models are trained on string molecular representations in Simplified Molecular Input Line Entry System (SMILES) [28] format to learn latent representations.

4.1. Text encoders

We used the following versions of BERT-based [1]:

- For English data, we used the large version of RoBERTa [18], pretrained on English texts of various domains. The total size of pretraining corpora is 160 GB. The model has 16 heads, 24 layers, 1024 embedding size, and a total of 355M parameters;
- For Russian data, we used EnRuDR-BERT,² a bilingual model pretrained on the (i) RuDReC [20] corpus, a Russian corpus of drug reviews, and (ii) the English corpus, the same one used for pretraining of EnDR-BERT. The model has 12 heads, 12 layers, 768 embedding size, and a total of 110M parameters;
- For French data, we used CamemBERT_{base} [29], a RoBERTa-based [18] model, pretrained on the French subcorpus of OSCAR corpus [30]. The model has 12 heads, 12 layers, 768 embedding size, and a total of 110M parameters;

4.2. Drug encoders

To obtain drug information, we map drug names to their identifiers in DrugBank [31] using the DrugBank terminology. The DrugBank database is a unique bioinformatics and cheminformatics resource that combines detailed drug (i.e., chemical, pharmacological, and pharmaceutical) data with comprehensive drug target information. We use Drugbank v. 5.1.8 that contains 14 325 identifiers.

4.2.1. ATC classification

Nowadays, the Anatomical Therapeutic Chemical (ATC) classification system is the most widely recognized classification system for drugs. The World Health Organization (WHO) recommended the ATC classification system for the division of drugs into diverse groups according to the organ or system on which they act and/or their therapeutic and chemical characteristics. This system classifies drugs into 14 main groups³: (A) alimentary tract and metabolism; (B) blood and blood-forming organs; (C) cardiovascular system; (D) dermatologicals; (G) genitourinary system and sex hormones; (H) systemic hormonal preparations, excluding sex hormones and insulins; (J) anti-infectives for systemic use; (L) antineoplastic and immunomodulating agents; (M) musculoskeletal system; (N) nervous system; (P) antiparasitic products, insecticides, and repellents; (R) respiratory system; (S) sensory organs; (V) various. [32] used chemical–chemical interactions and similarities to predict ATC codes of drugs since compounds with similar physicochemical properties are often involved in similar biological activities. Following this assumption, we encode a drug as a sparse 14-dimensional vector with zero values for ATC codes to which a drug does not belong. We adopt information about ATC main groups from Drugbank for each drug.

4.2.2. Molecular descriptors

Similar molecular structures have similar molecular properties. A molecular descriptor is defined as the “final result of a logical and mathematical procedure, which transforms chemical information encoded within a symbolic representation of a molecule into a useful number or the result of some standardized experiment” [33]. We use molecular-descriptor calculation library Mordred [16] to calculate 2 thousand descriptors for each drug from Drugbank.

4.2.3. MolBERT

MolBERT uses two domain-specific training objectives to improve on BERT: SMILES equivalence and calculated molecular descriptor

² <https://huggingface.co/cimm-kzn/enrudr-bert>.

³ http://www.whocc.no/atc/structure_and_principles/.




Fig. 2. Percentage of tweets containing ADE mentions grouped by ATC group for different SMM4H corpora. ATC groups with no positive class samples are not shown.

prediction [14]. They also used traditional masked language modeling. The SMILES equivalence task is to predict whether the two SMILES represent the same molecule based on the cross-entropy loss. The descriptor prediction task is to predict the set of descriptors for each molecule based on the mean squared error over all predicted values and scores from RDKit. The final loss is given by the arithmetic mean of all task losses. The authors used 1.5M compounds for pre-training. We use the official checkpoint with 12 attention heads and 12 layers

trained for 100 epochs and available at <https://github.com/BenevolentAI/MolBERT>.

4.2.4. ChemBERTa

ChemBERTa follows the RoBERTa architecture with 12 attention heads and 6 layers [15]. The authors adopted a pretraining procedure from RoBERTa, which masks 15% of the tokens in each input string. They used a dataset of 77M unique SMILES for pre-training. The official

checkpoint are available at https://huggingface.co/seyonec/ChemBERTa_zinc250k_v2_40k.

4.3. Classification network

For classification, we utilize a fully-connected classification network. The classifier includes one hidden layer, GeLU [34] activations and 0.3% dropout. We use sigmoid as an output layer activation and binary cross-entropy as the loss function. We freeze the first 5 Transformer encoder layers and the embedding layer of the encoder that shows higher results on the dev set.

5. Experiments

5.1. Experimental setup

We compared the performance of our proposed bi-modal models with three uni-modal (text-only) baselines:

- Support Vector Machines (SVM) [35] classifier over TF-IDF vectors. Before TF-IDF vectorization, we tokenized and lemmatized each text using Spacy⁴ and filtered out stop words using language-specific lists of stopwords obtained from NLTK [36]. We used TF-IDF and SVM implementations from Scikit-learn [37].
- Convolutional Neural Network (CNN) [38,39] over pretrained Word2vec [40] or Fasttext [41] word embeddings. We tokenized each text using Spacy and filtered out stop words using stopwords lists obtained from NLTK. Domain-specific biomedical English Word2Vec⁵ and Russian FastText⁶ models and a general-domain French FastText⁷ model were used to obtain word embeddings for the CNN baseline. Each CNN model was composed of three parallel CNN layers followed by Rectified Linear Unit (ReLU) activation and one-dimensional max-pooling. A final text representation was obtained as a concatenation of these three resulting representations fed to a fully-connected layer. Each of the three CNN layer had 128 filters with window sizes of 3, 4, and 5. We trained each CNN model for 15 epochs with a learning rate of $1 \cdot 10^{-4}$ and made a prediction using model parameters from the best epoch in terms of validation F1-score.
- BERT-based unimodal classifier.

For Russian and French datasets, we trained each BERT-based model for 10 epochs with a batch size of 64 and a learning rate of $3 \cdot 10^{-5}$ using Adam optimizer [42]. For test set predictions, we used model weights from the best epoch in terms of dev set F1-score. For English data, we trained each RoBERTa-based model for 5 epochs with the learning rate of $1 \cdot 10^{-5}$ and used the model weights from the last training epoch to make a text set prediction. Since English and French training sets have a strong class imbalance, we applied a positive class over-sampling to English and French corpora.

CNN baselines, unimodal and bi-modal BERT-based models were implemented using the PyTorch library [43]. Pretrained BERT-based models were obtained using the Transformers library [44]. For each dataset and training setup, we trained 10 BERT-based models with different initializations of classifier weights and calculated the mean quality and standard deviation of F₁-scores (F₁), precision (P), and recall (R) for the ADE class.

Table 3

Evaluation results for SMM4H 2020 Task 1b on classification of French tweets.

Model	P ± std	R ± std	F ₁ ± std
Text-only models			
TF-IDF + SVM	0.769	0.2	0.111
Fasttext + CNN	0.018	0.7	0.034
BERT	0.196 ± 0.170	0.2 ± 0.133	0.169 ± 0.108
[19]	0.15	0.20	0.17
BERT	0.196 ± 0.170	0.2 ± 0.133	0.169 ± 0.108
Models with concatenated modalities			
BERT + ATC categories	0.185 ± 0.049	0.23 ± 0.134	0.190 ± 0.063
BERT + descriptors	0.317 ± 0.176	0.22 ± 0.092	0.245 ± 0.111*
BERT + ChemBERTa	0.204 ± 0.221	0.14 ± 0.135	0.154 ± 0.151
BERT + MolBERT	0.415 ± 0.233	0.23 ± 0.106	0.254 ± 0.071*
Models with cross-attention			
BERT + ATC categories	0.249 ± 0.302	0.09 ± 0.057	0.113 ± 0.069
BERT + descriptors	0.260 ± 0.199	0.14 ± 0.084	0.164 ± 0.086
BERT + ChemBERTa	0.225 ± 0.083	0.17 ± 0.082	0.187 ± 0.075
BERT + MolBERT	0.251 ± 0.142	0.16 ± 0.117	0.178 ± 0.099

*Denotes statistically significantly better than unimodal BERT (Wilcoxon, $\rho < 0.01$).

Table 4

Evaluation results for SMM4H 2021 Task 2 on classification of Russian tweets [3]. We did not observe a statistically significant (Wilcoxon, $\rho < 0.05$) difference between bi-modal models and unimodal BERT.

Model	P ± std	R ± std	F ₁ ± std
Text-only models			
TF-IDF + SVM	0.158	0.281	0.202
Fasttext + CNN	0.356	0.465	0.404
BERT	0.548 ± 0.063	0.516 ± 0.072	0.524 ± 0.020
Models with concatenated modalities			
BERT + ATC categories	0.529 ± 0.054	0.521 ± 0.058	0.519 ± 0.090
BERT + descriptors	0.543 ± 0.052	0.514 ± 0.059	0.523 ± 0.020
BERT + ChemBERTa	0.552 ± 0.061	0.513 ± 0.076	0.524 ± 0.023
BERT + MolBERT	0.538 ± 0.040	0.531 ± 0.040	0.532 ± 0.014
Models with cross-attention			
BERT + ATC categories	0.509 ± 0.063	0.542 ± 0.059	0.519 ± 0.011
BERT + descriptors	0.527 ± 0.076	0.518 ± 0.073	0.514 ± 0.022
BERT + ChemBERTa	0.515 ± 0.049	0.537 ± 0.063	0.521 ± 0.020
BERT + MolBERT	0.514 ± 0.051	0.553 ± 0.055	0.529 ± 0.011

5.2. ADE text classification

Table 3 presents the results on the French set for SMM4H 2020 Task 1b. Based on the results, we can draw the following conclusions. First, the highest F₁-scores were achieved by bi-modal models with concatenated modalities. In particular, the bi-modal model with MolBERT drug embedding outperformed the best official results of the SMM4H 2020 competition and text-only model by 8.5% in terms of F₁-score, achieving a new state-of-the-art. Second, the performance of models with cross-attention is significantly lower than the ones with modalities concatenation. However, cross-attention models with chemBERTa and MolBERT embeddings slightly outperformed the text-only model by 1.8% and 0.9%. Finally, uni-modal and bi-modal BERT-based models significantly outperform unimodal SVM and CNN baselines, thus proving the superiority of complex neural networks over simple neural models and traditional machine learning algorithms in the task of ADE text classification.

Table 4 shows the performance of uni-modal and bi-modal models on the Russian tweet corpus of SMM4H 2021. According to the results, we can make the following observations. First, the two best-performing models in terms of F₁-score are bi-modal models that utilize MolBERT drug embeddings. Concatenation-based model slightly outperforms attention-based model (0.3%) and unimodal BERT (0.8%). Second, bi-modal models with concatenated modalities outperform the ones with cross-attention. Third, our final SMM4H 2021 submission combined the results of ten multimodal models with the same

⁴ <https://spacy.io>.

⁵ <https://bio.nplab.org>.

⁶ <https://github.com/cimm-kzn/RuDReC>.

⁷ <https://fasttext.cc/docs/en/crawl-vectors.html>.

Table 5

Evaluation results for on classification of English tweets on SMM4H 2021 dataset (unofficial class-stratified test set).

Model	P \pm std	R \pm std	F ₁ \pm std
Text-only models			
TF-IDF + SVM	0.484	0.417	0.448
Word2vec + CNN	0.664	0.344	0.453
RoBERTa	0.755 \pm 0.062	0.842 \pm 0.051	0.793 \pm 0.022
Team 4 [45]	0.59	0.92	0.72
Models with concatenated modalities			
RoBERTa + ATC categories	0.761 \pm 0.048	0.829 \pm 0.064	0.791 \pm 0.025
RoBERTa + descriptors	0.737 \pm 0.045	0.838 \pm 0.051	0.782 \pm 0.016
RoBERTa + ChemBERTa	0.757 \pm 0.055	0.828 \pm 0.057	0.787 \pm 0.014
RoBERTa + MolBERT	0.762 \pm 0.044	0.829 \pm 0.057	0.792 \pm 0.019
Models with cross-attentions			
RoBERTa + ATC categories	0.770 \pm 0.051	0.812 \pm 0.058	0.788 \pm 0.022
RoBERTa + descriptors	0.794 \pm 0.031	0.808 \pm 0.047	0.799 \pm 0.021*
RoBERTa + ChemBERTa	0.771 \pm 0.065	0.822 \pm 0.054	0.792 \pm 0.025
RoBERTa + MolBERT	0.760 \pm 0.057	0.832 \pm 0.034	0.793 \pm 0.031

*Denotes statistically significantly better than unimodal RoBERTa (Wilcoxon, $p < 0.01$).

Table 6

Official SMM4H 2021 results [3] on the official test set.

Model	P	R	F ₁
Task 2, Russian			
Our SMM4H 2021 Task 2 submission [17]	0.58	0.57	0.57
Second place [3]	0.54	0.57	0.52
Task 1a, English			
Our SMM4H 2021 Task 1a submission [17]	0.552	0.681	0.61
First/second place, Team 4 [45]	0.515	0.752	0.61

Table 7

Evaluation results for classification of English tweets on the ATC-stratified SMM4H 2021 test set.

Model	P \pm std	R \pm std	F ₁ \pm std
Text-only models			
RoBERTa	0.682 \pm 0.058	0.764 \pm 0.076	0.715 \pm 0.032
Models with concatenated modalities			
RoBERTa + ATC categories	0.675 \pm 0.040	0.782 \pm 0.038	0.723 \pm 0.019
RoBERTa + descriptors	0.679 \pm 0.059	0.786 \pm 0.044	0.725 \pm 0.019*
RoBERTa + ChemBERTa	0.685 \pm 0.049	0.778 \pm 0.037	0.726 \pm 0.014
RoBERTa + MolBERT	0.654 \pm 0.055	0.792 \pm 0.051	0.713 \pm 0.018
Models with cross-attentions			
RoBERTa + ATC categories	0.682 \pm 0.063	0.762 \pm 0.057	0.716 \pm 0.022
RoBERTa + descriptors	0.684 \pm 0.062	0.762 \pm 0.051	0.717 \pm 0.017
RoBERTa + ChemBERTa	0.676 \pm 0.061	0.782 \pm 0.055	0.721 \pm 0.020
RoBERTa + MolBERT	0.668 \pm 0.070	0.784 \pm 0.053	0.717 \pm 0.030

*Denotes statistically significantly better than unimodal RoBERTa (Wilcoxon, $p < 0.05$).

settings using a simple voting scheme with the intent of increasing the robustness of the final system. As shown in Table 6, this model achieved state-of-the-art results showing 4% improvement in terms of F₁-measure over a single model.

Table 5 shows the performance of uni-modal and bi-modal models on English tweets. Since there are strict limits with 3 submissions per day for the official test set on the SMM4H 2021 Codalab page,⁸ we use 20% of the original train set as a test set for the evaluation of our models. Several observations can be made based on Table 5. First, for the dev set, the majority of bi-modal models do not show a significant performance difference compared to the uni-modal classifier, but the highest F₁-score is achieved by the descriptors-based model with cross-attention that outperforms the text-only model by 0.6% in terms of F₁.

Table 8

Comparison of uni- and bi-modal models performance on various therapeutic groups. The number in parenthesis indicates the number of tweets about drugs associated with the ATC class.

ATC group	Uni-modal model			Bi-modal model		
	P	R	F ₁	P	R	F ₁
SMM4H 2021 Task 2, Russian tweets, official test set						
Nervous system (6100)	0.456	0.562	0.495	0.472	0.558	0.506
Anti-infectives for systemic use (2859)	0.440	0.505	0.464	0.448	0.495	0.465
Sensory organs (365)	0.767	0.720	0.738	0.805	0.747	0.771
Alimentary tract and metabolism (280)	0.675	0.669	0.662	0.713	0.652	0.670
Dermatologicals (61)	0.967	0.725	0.825	0.974	0.75	0.845
SMM4H 2021 Task 1a, English tweets, official dev set						
Nervous system (470)	0.734	0.864	0.791	0.785	0.830	0.805
Dermatologicals (158)	0.613	0.842	0.703	0.626	0.808	0.705
Alimentary tract and metabolism (155)	0.767	0.85	0.804	0.816	0.75	0.774
Respiratory system (130)	0.658	0.925	0.760	0.678	0.875	0.760
Sensory organs (88)	0.832	0.880	0.851	0.849	0.79	0.814

Second, the best F₁ on the our test set is significantly higher than on the official test set (0.799 vs. 0.61). Third, as shown in Table 6, our official SMM4H 2021 submission combined the results of ten models with the same settings using a simple voting scheme achieved the results on par with the best-performing team [45] that utilized the RoBERTa model with undersampling and oversampling.

5.2.1. Performance on ATC-stratified test set

In this subsection, we investigate if bi-modal models perform better in a case when train and test sets are similar in terms of covered ATC classes. For this purpose, we compared the proposed bi-modal models and unimodal BERT on ATC-stratified SMM4H 2021 English data split, i.e., on a dataset that has the same distribution of ATC groups among train and test sets. Table 7 presents the results. The following observations can be made. Bi-modal models that are based on ATC groups, ChemBERTa embeddings, and chemical descriptors consistently perform no worse than the unimodal BERT in terms of F₁-score and have lower performance variance. The highest quality is observed for concatenation-based ChemBERTa and descriptors models that outperform unimodal BERT by 1.1% and 1.0% F₁, respectively. But the difference between bi-modal and unimodal models is statistically significant ($p \leq 0.05$) for the descriptors-based model with modalities concatenation only. Based on the results, we cannot conclude there exists a dependency between the similarity of train and test sets' drug groups distribution and the quality of a bi-modal model.

5.2.2. Performance on different drug groups

In this subsection, we explore how the incorporation of additional modality affects the classification performance on tweets that mention drugs of different therapeutic groups. For the Russian and English corpus of SMM4H 2021, we took top-performing bi-modal models in terms of F₁-score and compared their performance with uni-modal models on test set subsets of 5 most frequent ATC groups.

Table 8 presents the results for the English and Russian corpora of SMM4H 2021. Based on the results, we can draw the following conclusions. First, bi-modal models show an increase of more than 1% F₁ compared to uni-modal models for the most frequent ATC group "Nervous system" on both datasets. Second, for each dataset and therapeutic group, bi-modal models obtain higher precision. The precision growth varies from 0.7% ("Dermatologicals") to 3.8% ("Sensory organs") for the Russian corpus and from 1.3% ("Dermatologicals") to 5.1% ("Nervous system") for English corpus. Third, for all top-5 therapeutic groups of Russian corpus, bi-modal models achieved higher F₁. The highest F₁ increase is achieved for group "Sensory organs" (+3.3%), and the lowest increase of 0.1% is insignificant and is achieved for group "Anti-infectives for systemic use". For English data, the use of bi-modal results in F₁ decrease for groups "Alimentary tract and metabolism"

⁸ <https://competitions.codalab.org/competitions/28766>.

Table 9

SMM4H 2021 data samples with predictions by our model. The “Error type” column presents prediction error type according to the human annotation: “False positive” stands for the cases when the model incorrectly labels a text sample as ADE, and “False negative” stands for the cases when the model incorrectly labels a text sample as non-ADE.

Error reason	Text sample	Error type
Positive effect	However the prozac is making her more talkative than usual (if also hilarious and lucid) so i don't think she's fully listening to my ideas	False positive
Positive effect	Nicotine always makes me get out of a bad mood	False positive
Positive effect	Hyper af! I feel like a bird on vyvanse	False positive
Positive effect	не знаю от чего меня так выптырило от ношпы или парацетамола но я не могу перестать танцевать и мне в общем то хорошо (I don't know if Drotaverine or Paracetamol got me this feeling, but I feel pretty good and can't stop dancing)	False positive
Annotation error	Yay! Cymbalta is making me numb again.	False positive
Annotation error	Alright time to find out if a ton of vyvanse and not sleeping for 2 days can get a trimester worth of homework done for 2 classes	False negative
Ambiguous sample	Half a trazodone had me sleeping heavy this morning....	False positive
Ambiguous sample	Последствия наложения парацетамола на моё организм: проснулся с температурой 35. (The consequences of applying paracetamol on my body: I woke up with a temperature of 35.0)	False positive

(−3.0%) and “Sensory organs” (−3.7%). Finally, for a fixed therapeutic group, the sign of performance difference between bi-modal and uni-modal models varies for different languages. In particular, for group “Sensory organs” the bi-modal model shows a 3.3% growth and 3.7% decrease of F_1 for Russian and English corpus, respectively. Hence, we can conclude that the performance of the bi-modal classification approach is not dependent on only ATC groups’ distribution of input data.

5.3. Error analysis

Table 9 shows error analysis for sample texts predicted by our model. First, the samples that mention an unexpected positive drug effect are often wrongly labeled as ADE samples. It seems like the modern BERT-based models fail to distinguish the beneficial drug effects from the unpleasant ones. Second, the SMM4H corpora are noisy and contain annotation errors. Finally, there are ambiguous samples that mention a drug effect but do not specify if it is a desired effect or an adverse one. The existing text-only classification models are bound to fail on such samples as the textual modality does not provide enough information for proper classification. Thus, the models have to utilize external knowledge that can be presented in drug metadata explicitly or implicitly encoded in a drug representation.

6. Discussion

In our work, we propose a multimodal model based on text and drug representations. Our experiments on three datasets demonstrate that the proposed model outperforms models based on text representations only. We have shown how our model improves upon the current state of the art and analyzed the influence of training set size on the results.

First, we investigated monolingual models only. These models still require manually annotated training data, which is not available for many languages. We can investigate how to transfer knowledge from one language to another. We believe that drug representations can help to gain improvement in this setup, since multilingual representations

are not perfectly aligned, especially, for domain-specific topics such as biomedicine. In contrast, drug representations will be the same and can serve as implicit knowledge anchors about differences in drug groups.

Second, we can encode additional knowledge into drug representations by training multi-label classifiers for the prediction of the side effects of a given drug on the SIDER dataset. Another direction is multimodal multitask learning on textual datasets and the SIDER dataset that will require jointly modeling the drug and the text modalities. We expect that further improvements may be possible with more research investigating other networks and developing a better joint loss function for two tasks.

Third, zero-shot cross-drug class evaluation is still an open research direction. A recurring limitation, which arises with supervised models, is bias in training data; annotated corpora include texts about a small number of drugs (at best about dozens) while there are over 20,000 prescription drug products approved for marketing.⁹ The model performance is evaluated under the implicit hypothesis that the training data (source) and the test data (target) come from the same underlying distribution (i.e., both sets include drugs from the same ATC main group). It would be interesting to rebuild train/dev/test sets to model cross-drug class evaluation.

7. Conclusion

In this paper, we studied the task of discovering adverse drug effects (ADEs) presence in user-generated tweets about drugs. In this work, we have presented a unified model that combines several state-of-the-art models, in particular BERT for text representation and MolBERT for drug representation. Our models led to significant improvements on ADE classification of French texts (by 8% with the previous state-of-the-art quality of 17% F1) and achieved state-of-the-art results on recent SMM4H 2021 Task 1a and Task 2 for English and Russian texts, respectively, all of this using only well-known BERT-based pre-trained models for individual components. Possible directions for future work include: (i) multilingual learning on several datasets; (ii) multi-task

⁹ <https://www.fda.gov/about-fda/fda-basics/fact-sheet-fda-glance>.

learning on two tasks: ADE text classification and prediction of the side effects of a given drug using drug structure or knowledge graphs with interactions between biomedical entities; (iii) zero-shot cross-drug class evaluation across drug classes.

CRedit authorship contribution statement

Andrey Sakhovskiy: Implemented the machine learning framework, Conducted experiments, Error analysis, Manuscript writing. **Elena Tutubalina:** Methodology, Collected the datasets, Manuscript writing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The work on novel multimodal models and manuscript has been supported by the Russian Science Foundation grant # 18-11-00284. The work on baselines has been supported by a grant from the President of the Russian Federation for young scientists-candidates of science (MK-3193.2021.1.6).

References

- [1] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in: 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), 2019, pp. 4171–4186.
- [2] A. Klein, I. Alimova, I. Flores, A. Magge, Z. Miftahutdinov, A.-L. Minard, K. O'Connor, A. Sarker, E. Tutubalina, D. Weissenbacher, et al., Overview of the fifth social media mining for health applications (# smm4h) shared tasks at coling 2020, in: Proceedings of the Fifth Social Media Mining for Health Applications Workshop & Shared Task, 2020, pp. 27–36.
- [3] A. Magge, A. Klein, A. Miranda-Escalada, M.A. Al-Garadi, I. Alimova, Z. Miftahutdinov, E. Farre, S. Lima-López, I. Flores, K. O'Connor, et al., Overview of the sixth social media mining for health applications (# smm4h) shared tasks at naacl 2021, in: Proceedings of the Sixth Social Media Mining for Health (# SMM4H) Workshop and Shared Task, 2021, pp. 21–32.
- [4] Z. Miftahutdinov, A. Sakhovskiy, E. Tutubalina, KFU NLP team at SMM4H 2020 tasks: Cross-lingual transfer learning with pretrained language models for drug reactions, in: Proceedings of the Fifth Social Media Mining for Health Applications Workshop & Shared Task, 2020, pp. 51–56.
- [5] Z. Miftahutdinov, I. Alimova, E. Tutubalina, On biomedical named entity recognition: experiments in interlingual transfer for clinical and social media texts, in: European Conference on Information Retrieval, Springer, 2020, pp. 281–288.
- [6] C.-K. Wang, H.-J. Dai, Y.-C. Zhang, B.-C. Xu, B.-H. Wang, Y.-N. Xu, P.-H. Chen, C.-H. Lee, ISLab system for SMM4H shared task 2020, in: Proceedings of the Fifth Social Media Mining for Health Applications Workshop & Shared Task, 2020, pp. 42–45.
- [7] A. Gusev, A. Kuznetsova, A. Polyanskaya, E. Yatsishin, BERT implementation for detecting adverse drug effects mentions in Russian, in: Proceedings of the Fifth Social Media Mining for Health Applications Workshop & Shared Task, 2020, pp. 46–50.
- [8] G.M. Dimitri, P. Lió, DrugClust: A machine learning approach for drugs side effects prediction, *Comput. Biol. Chem.* 68 (2017) 204–210.
- [9] S. Honda, S. Shi, H.R. Ueda, SMILES transformer: pre-trained molecular fingerprint for low data drug discovery, 2019, arXiv preprint [arXiv:1911.04738](https://arxiv.org/abs/1911.04738).
- [10] P. Li, Y. Li, C.-Y. Hsieh, S. Zhang, X. Liu, H. Liu, S. Song, X. Yao, TrimNet: learning molecular representation from triplet messages for biomedicine, *Brief. Bioinform.* (2020).
- [11] S. Wang, Y. Guo, Y. Wang, H. Sun, J. Huang, SMILES-BERT: large scale unsupervised pre-training for molecular property prediction, in: Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics, 2019, pp. 429–436.
- [12] M. Kuhn, I. Letunic, L.J. Jensen, P. Bork, The SIDER database of drugs and side effects, *Nucleic Acids Res.* 44 (D1) (2016) D1075–D1079.
- [13] Z. Wu, B. Ramsundar, E.N. Feinberg, J. Gomes, C. Geniesse, A.S. Pappu, K. Leswing, V. Pande, MoleculeNet: a benchmark for molecular machine learning, *Chem. Sci.* 9 (2) (2018) 513–530.
- [14] B. Fabian, T. Edlich, H. Gaspar, M. Segler, J. Meyers, M. Fiscato, M. Ahmed, Molecular representation learning with language models and domain-relevant auxiliary tasks, 2020, arXiv preprint [arXiv:2011.13230](https://arxiv.org/abs/2011.13230).
- [15] S. Chithrananda, G. Grand, B. Ramsundar, ChemBERTa: Large-scale self-supervised pretraining for molecular property prediction, 2020, arXiv preprint [arXiv:2010.09885](https://arxiv.org/abs/2010.09885).
- [16] H. Moriwaki, Y.-S. Tian, N. Kawashita, T. Takagi, Mordred: a molecular descriptor calculator, *J. Cheminform.* 10 (1) (2018) 1–14.
- [17] A. Sakhovskiy, Z. Miftahutdinov, E. Tutubalina, KFU NLP team at SMM4H 2021 tasks: Cross-lingual and cross-modal BERT-based models for adverse drug effects, in: Proceedings of the Sixth Social Media Mining for Health (# SMM4H) Workshop and Shared Task, 2021, pp. 39–43.
- [18] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, 2019, arXiv preprint [arXiv:1907.11692](https://arxiv.org/abs/1907.11692).
- [19] O. Gencoglu, Sentence transformers and Bayesian optimization for adverse drug effect detection from Twitter, in: Proceedings of the Fifth Social Media Mining for Health Applications Workshop & Shared Task, 2020, pp. 161–164.
- [20] E. Tutubalina, I. Alimova, Z. Miftahutdinov, A. Sakhovskiy, V. Malykh, S. Nikolenko, The Russian Drug Reaction Corpus and neural models for drug reactions and effectiveness detection in user reviews, *Bioinformatics* 37 (2) (2021) 243–249.
- [21] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, in: Proceedings of the 31st International Conference on Neural Information Processing Systems, 2017, pp. 6000–6010.
- [22] M. Asada, M. Miwa, Y. Sasaki, Enhancing drug-drug interaction extraction from texts by molecular structure information, in: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), 2018, pp. 680–685.
- [23] S. Saha, P. Dutta, Amalgamation of protein sequence, structure and textual information for improving protein-protein interaction identification, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 2020, pp. 6396–6407.
- [24] J. Yu, J. Li, Z. Yu, Q. Huang, Multimodal transformer with multi-view visual representation for image captioning, *IEEE Trans. Circuits Syst. Video Technol.* 30 (12) (2019) 4467–4480.
- [25] A. Savchenko, A. Alekseev, S. Kwon, E. Tutubalina, E. Myasnikov, S. Nikolenko, Ad lingua: Text classification improves symbolism prediction in image advertisements, in: Proceedings of the 28th International Conference on Computational Linguistics, 2020, pp. 1886–1892.
- [26] U.S. Food and Drug Administration, Drugs@FDA glossary of terms, 2017, <https://www.fda.gov/drugs/drug-approvals-and-databases/drugsfda-glossary-terms>, [Drug; Drug Product; online, accessed 21-July-2020].
- [27] Ministry of Health of the Russian Federation, State register of medicinal remedies, 2021, <https://grls.rosminzdrav.ru/Default.aspx>, [Drug; Drug Product; online, accessed 28-June-2021].
- [28] D. Weininger, SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules, *J. Chem. Inf. Comput. Sci.* 28 (1) (1988) 31–36.
- [29] L. Martin, B. Muller, P.J.O. Suárez, Y. Dupont, L. Romary, É.V. de la Clergerie, D. Seddah, B. Sagot, CamBERT: a tasty French language model, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 2020, pp. 7203–7219.
- [30] P.J.O. Suárez, B. Sagot, L. Romary, Asynchronous pipeline for processing huge corpora on medium to low resource infrastructures, in: 7th Workshop on the Challenges in the Management of Large Corpora (CMLC-7), Leibniz-Institut für Deutsche Sprache, 2019.
- [31] D.S. Wishart, C. Knox, A.C. Guo, D. Cheng, S. Shrivastava, D. Tzur, B. Gautam, M. Hassanali, DrugBank: a knowledgebase for drugs, drug actions and drug targets, *Nucleic Acids Res.* 36 (suppl_1) (2008) D901–D906.
- [32] L. Chen, W.-M. Zeng, Y.-D. Cai, K.-Y. Feng, K.-C. Chou, Predicting anatomical therapeutic chemical (ATC) classification of drugs by integrating chemical-chemical interactions and similarities, *PLoS One* 7 (4) (2012) e35254.
- [33] V. Consonni, R. Todeschini, Molecular descriptors, in: Recent Advances in QSAR Studies, Springer, 2010, pp. 29–102.
- [34] D. Hendrycks, K. Gimpel, Gaussian error linear units (gelus), 2016, arXiv preprint [arXiv:1606.08415](https://arxiv.org/abs/1606.08415).
- [35] B.E. Boser, I.M. Guyon, V.N. Vapnik, A training algorithm for optimal margin classifiers, in: Proceedings of the Fifth Annual Workshop on Computational Learning Theory, 1992, pp. 144–152.
- [36] S. Bird, E. Klein, E. Loper, Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit, " O'Reilly Media, Inc.", 2009.

- [37] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.
- [38] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* 86 (11) (1998) 2278–2324.
- [39] Y. Kim, Convolutional neural networks for sentence classification, in: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP, 2014*, pp. 1746–1751.
- [40] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, in: *Proceedings of Workshop at ICLR, 2013*, 2013.
- [41] P. Bojanowski, E. Grave, A. Joulin, T. Mikolov, Enriching word vectors with subword information, *Trans. Assoc. Comput. Linguist.* 5 (2017) 135–146.
- [42] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Y. Bengio, Y. LeCun (Eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015*, URL <http://arxiv.org/abs/1412.6980>.
- [43] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, PyTorch: An imperative style, high-performance deep learning library, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32*, Curran Associates, Inc., 2019, pp. 8024–8035, URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [44] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T.L. Scao, S. Gugger, M. Drame, Q. Lhoest, A.M. Rush, Transformers: State-of-the-art natural language processing, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Association for Computational Linguistics, Online, 2020, pp. 38–45, URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- [45] S. Ramesh, A. Tiwari, P. Choubey, S. Kashyap, S. Khose, K. Lakara, N. Singh, U. Verma, BERT based transformers lead the way in extraction of health information from social media, in: *Proceedings of the Sixth Social Media Mining for Health (# SMM4H) Workshop and Shared Task, 2021*, pp. 33–38.