

Methods for Detoxification of Texts for the Russian Language

Daryna Dementieva[‡], Daniil Moskovskiy[‡], Varvara Logacheva[‡], David Dale[‡],
Olga Kozlova[†], Nikita Semenov[†], and Alexander Panchenko[‡]

[‡]Skolkovo Institute of Science and Technology, Moscow, Russia

[†]Mobile TeleSystems (MTS), Moscow, Russia

{daryna.dementieva, daniil.moskovskiy, v.logacheva, d.dale, a.panchenko}@skoltech.ru

{oskozlo9,nikita.semenov}@mts.ru

Abstract

We introduce the first study of automatic detoxification of Russian texts to combat offensive language. Such a kind of textual style transfer can be used, for instance, for processing toxic content in social media. While much work has been done for the English language in this field, it has never been solved for the Russian language yet. We test two types of models – unsupervised approach based on BERT architecture that performs local corrections and supervised approach based on pretrained language GPT-2 model – and compare them with several baselines. In addition, we describe evaluation setup providing training datasets and metrics for automatic evaluation. The results show that the tested approaches can be successfully used for detoxification, although there is room for improvement.

Keywords: text style transfer, toxicity detection, detoxification, pre-trained models

DOI: 10.28995/2075-7182-2021-20-179-190

Методы детоксификации текстов для русского языка

Дарина Дементьева[‡], Даниил Московский[‡], Варвара Логачева[‡], Давид Далё[‡],
Ольга Козлова[†], Никита Семенов[†], Александр Панченко[‡]

[‡]Сколковский Институт Науки и Технологий, Москва, Россия

[†]Мобильные ТелеСистемы (МТС), Москва, Россия

{daryna.dementieva, daniil.moskovskiy, v.logacheva, d.dale, a.panchenko}@skoltech.ru

{oskozlo9,nikita.semenov}@mts.ru

Аннотация

Мы представляем первое в своем роде исследование автоматической детоксикации русскоязычных текстов для борьбы с оскорбительной речью. Такой перенос стиля для текстов может быть использован, например, для предварительной обработки в социальных сетях. В то время как решения подобных задач уже были представлены для английского языка, для русского такая постановка задачи и методы её решения описываются впервые. Мы провели эксперименты по тестированию двух типов моделей – метод обучения без учителя на основе архитектуры BERT, который выполняет локальные коррекции, и метод обучения с учителем на основе предобученной языковой модели GPT-2 – и сравнили их с несколькими базовыми подходами. Кроме того, мы предоставили описание методологии оценки вместе с набором обучающих данных и метрик для автоматической оценки. Результаты показали, что протестированные методы могут быть успешно использованы для детоксикации, однако могут быть усовершенствованы.

Ключевые слова: перенос стиля для текстов, определение токсичности, детоксификация, предобученные модели

1 Introduction

Global access to the Internet has enabled the spread of information all over the world and has given many new possibilities. On the other hand, alongside the advantages, the exponential and uncontrolled growth of user-generated content on the Internet has also facilitated the spread of toxicity and hate speech. Much work has been done in the direction of offensive speech detection [5, 23, 17]. However, it has become

essential not only to detect toxic content but also to combat it in smarter ways. While some social networks block sensitive content, another solution can be to detect toxicity in a text which is being typed in and offer a user a non-offensive version of this text. This task can be considered a style transfer task, where the source style is toxic, and the target style is neutral/non-toxic.

The task of style transfer is the task of transforming a text so that its content and the majority of properties stay the same, and one particular attribute (*style*) changes. This attribute can be the sentiment [24, 15], the presence of bias [19], the degree of formality [22], etc. The work [7] gives more examples of style transfer applications. Considering the task of detoxification, it has already been tackled by different groups of researchers [16, 26], as well as a similar task of transforming text to a more polite form [13]. However, all these works deal only with the English language. As for Russian, the methods of text style transfer and text detoxification have not been explored before.

To the best of our knowledge, our work is the first effort to solve the text style transfer task with a focus on toxicity elimination for the Russian language. We leverage pre-trained language models (GPT and BERT) and demonstrate that they can successfully solve the task after being trained on a very small parallel corpus or only on non-parallel data.

The contributions of this work are three-fold:

1. We introduce the new study of text detoxification for the Russian language;
2. We conduct experiments with two well-performing style transfer methods: a method based on GPT-2 which rewrites the text and a BERT-based model which performs targeted corrections;
3. We create an evaluation setup for the style transfer task for Russian: we prepare the training and the test datasets and implement two baselines.

2 Problem Statement

The definition of *textual style* in the context of NLP is still vague [25]. One of the first definitions of style refers to how the sense is expressed [14]. However, in our work, we adhere to the data-driven definition of style. Thus, the style simply refers to the characteristics of a given corpus that are distinct from a general text corpus [7]. The style is a particular characteristic from a set of categorical values: {positive, negative} [24], {polite, impolite} [13], {formal, informal} [22]. Commonly, it is assumed that this textual characteristic is measurable using a function $g(x_i) \rightarrow s_i$ that gets as input text x_i and returns the corresponding style label s_i . For instance, it can be implemented using a text classifier.

We define the task of style transfer as follows. Let us consider two corpora $D^X = \{x_1, x_2, \dots, x_n\}$ and $D^Y = \{y_1, y_2, \dots, y_m\}$ in two different styles – s^X and s^Y , respectively. The task is to create a model $f_\theta : X \rightarrow Y$, where X and Y are all possible texts with styles s^X and s^Y . The task of selecting the optimal set of parameters θ for f consists of maximising the probability $p(y'|x, s^Y)$ of transferring a sentence x with the style s^X to the sentence y' which saves the content of x and has the style s^Y . The parameters are maximised on the corpora D^X and D^Y which can be parallel or non-parallel. We focus on the transfer $s^X \rightarrow s^Y$, where s^X is the toxic style, and s^Y is neutral.

3 Related Work

Style transfer was first proposed and widely explored for images [6]. However, the task of text style transfer has currently gained less attention, partly due to the ambiguity of the term “style” for texts. Nevertheless, there exists a large body of work on textual style transfer for different styles. All the existing methods can be divided into techniques that use parallel training corpora and those using only non-parallel data. The latter category is larger because pairs of texts which share the content but have different styles are usually not available. At the same time, it is relatively easy to find non-parallel texts of the same domain with different styles (e.g. positive and negative movie reviews, speeches by politicians from different parties, etc.).

One of the methods which uses only non-parallel data is *Delete, Retrieve, Generate* [12] model. It is based on the idea that words in a sentence can be divided into those responsible for the sentence semantics and those carrying the style information. Therefore, if we delete the style words and replace them with

the corresponding words of the opposite style, we can change the style of the sentence while keeping the content intact. Alternative to this approach are methods that create disentangled representations of text [8]. In this case, the style and the content of a text are encoded into different spaces. When generating a text with a new style, we substitute the vector of the text style with the vector representation of the target style and generate a new sequence.

On the other hand, if there exists a corpus with parallel sentences $\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$ then style transfer can be formulated as a sequence-to-sequence task, analogously to supervised Machine Translation, summarization, paraphrasing, etc. Such models can greatly benefit from pre-trained language models, such as GPT [20] or T5 [21]. They often perform well on a range of NLP tasks with no fine-tuning. Moreover, when a small training dataset is available, their performance improves even further. For example, in [9] a GPT-based model was fine-tuned on an automatically generated parallel corpus to transfer between multiple styles. The recently released ruGPT3¹ model allows us to leverage big textual data for the detoxification task in Russian.

4 Methodology

We suggest several solutions to the text detoxification task. We test a method based on the GPT model, which uses parallel data and a BERT-based solution trained solely on non-parallel corpora. We also implement several baselines.

4.1 Baselines

Duplicate This is a naive baseline that amounts to performing no changes to the input sentence. It represents a lower bound of the performance of style transfer models, i.e. it helps us check that the models do not contaminate the original sentence.

Delete This method eliminates toxic words based on a predefined toxic words vocabulary. The idea is often used on television and other media: rude words are bleeped out or hidden with special characters (usually an asterisk). The main limitation of this method is vocabulary incompleteness: we cannot collect all the rude and toxic words. Moreover, new offensive words and phrases can appear in the language that can be also concatenated with different prefixes and suffixes. On the other hand, this method can preserve the content quite well, except for the cases when toxic words contain meaning that is essential for the understanding of the whole text.

Retrieve This method introduced in [12] is targeted at improving the accuracy of style transfer. For a given toxic sentence, we retrieve the most similar non-toxic text from a corpus of non-toxic samples. In this case, we get a safe sentence. However, the preservation of the content depends on the corpus size and is likely to be very low.

4.2 detoxGPT

GPT-2 [20] is a powerful language model which can be adapted to a wide range of NLP tasks using a very small task-specific dataset. Until recently, there were no such models for Russian. The AI Journey competition² released the ruGPT3 model capable of generating coherent and sensible texts in Russian. We suggest using it for style transfer via the following setups:

- **zero-shot**: the model is taken as is (with no fine-tuning). The input is a toxic sentence which we would like to detoxify prepended with the prefix “Перепарафразируй” (rus. *Paraphrase*) and followed with the suffix “>>>” to indicate the paraphrasing task. ruGPT3 has already been trained for this task, so this scenario is analogous to performing paraphrasing. The schematic pipeline of this setup is presented in Figure 1.
- **few-shot**: the model is taken as is. Unlike the previous scenario, we give a prefix consisting of a parallel dataset $\{(t_1^X, t_1^Y), \dots, (t_n^X, t_n^Y)\}$ of toxic and neutral sentences in the following form:

¹<https://github.com/sberbank-ai/ru-gpts>

²<https://ai-journey.ru>

“ $t_i^X \ggg t_i^Y$ ”. These examples can help the model understand that we require *detoxifying* paraphrasing. The parallel sentences are followed with the input sentence which we would like to detoxify with the prefix “Перепаразируй” and the suffix \ggg . The schematic pipeline of this setup is presented in Figure 2.

- ***fine-tuned***: the model is fine-tuned for the paraphrasing task on a parallel dataset $\{(t_1^X, t_1^Y), \dots, (t_n^X, t_n^Y)\}$. This implies training of the model on strings of the form “ $t_i^X \ggg t_i^Y$ ”. After the training, we give the input to the model analogously to the other scenarios. The schematic pipeline of this setup is presented in Figure 3.

zero-shot detoxGPT




Figure 1: The illustration of pipeline of *zero-shot* setup of detoxGPT approach.

few-shot detoxGPT




Figure 2: The illustration of pipeline of *few-shot* setup of detoxGPT approach.

fine-tuned detoxGPT

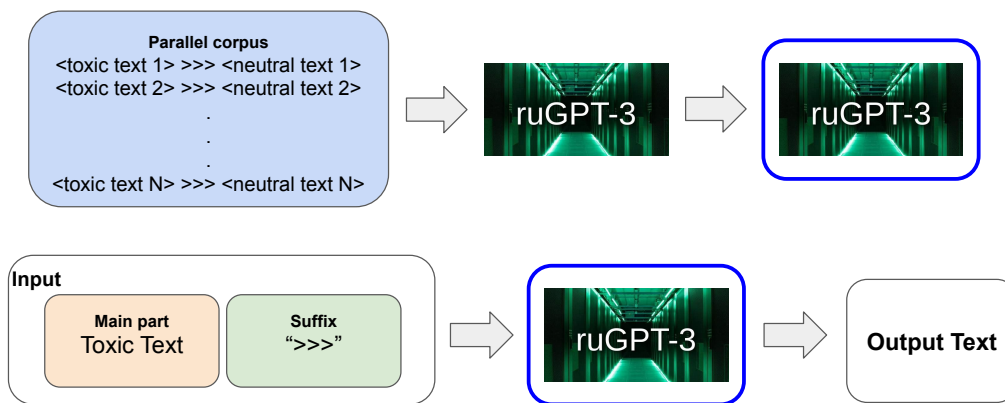


Figure 3: The illustration of pipeline of *fine-tuned* setup of detoxGPT approach.

The described methods require parallel data. These have to be pairs of sentences with the same content and the different toxicity level. Such sentences are not created “naturally” (unlike translations of the same text into different languages), so they have to be written from scratch to train such models. This is a laborious process. However, our intuition is that the detoxGPT model can perform detoxification after being trained on a very small number (several hundred) of parallel sentences, which can be created quickly.

4.3 condBERT

BERT (Bidirectional Encoder Representations from Transformers) [4] is a masked language model which has been trained on the task of predicting a missing word given the rest of the sentence. Although BERT is mainly used for getting word vector representations or sequence labeling and text classification tasks, it can also be used in the gap-filling scenario, i.e. for retrieving a word in a context that has been replaced with a [MASK] token. This scenario perfectly suits the delete-retrieve-generate style transfer method, which replaces individual words of a sentence and, as a result, generates so-called “lexical substitution” [2].

To make BERT fully suitable for style transfer, we need to change the model so that masking and replacing words changes the style of the input sentence. This can be done via fine-tuning BERT on style-specific corpora for the source and the target styles so that it learns the word distributions conditioned on a style and makes replacements that agree with it. Such a BERT-based model was first applied to the data augmentation task in [27]. Then, in [28], a similar model was used for sentiment style transfer.

The model **condBERT** (conditional BERT) model was proposed in [27]. While the tokens to replace were selected randomly in the original work, we mask tokens associated with the source style (toxic). To select the toxic words, we train a bag-of-words logistic regression model, which classifies the sentences as toxic or neutral. As a by-product of this model, we acquire weights for each word from the vocabulary. These weights can be interpreted as the toxicity level. We consider a token to be toxic if its weight is higher than a predefined threshold.

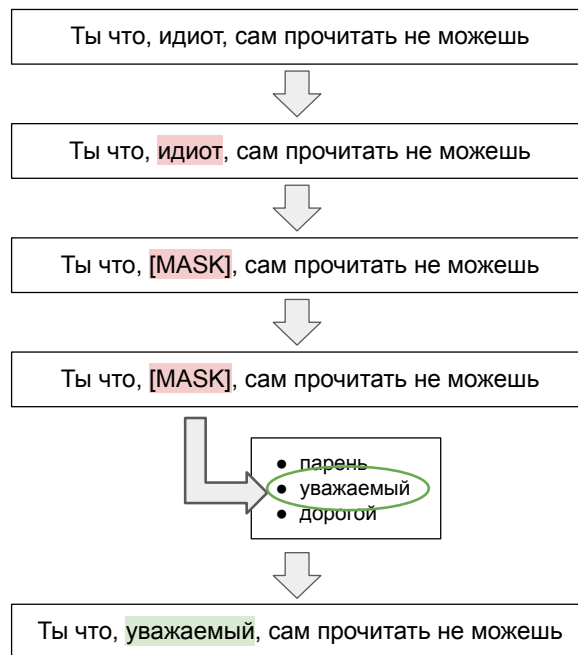


Figure 4: The illustration of the main idea of condBERT approach.

We then train the model on two corpora D^X and D^Y for the source and the target styles. To teach the model to distinguish styles, we include the style information as an extra embedding layer as described in [27]. Thus, it learns different distributions for toxic and non-toxic texts. To further force the model to replace toxic tokens with tokens that have a close meaning and are not toxic, we calculate the toxicity level of each token in the BERT vocabulary (using the logreg weights) and penalize the predicted probabilities of tokens that have a high toxicity. Finally, we enable condBERT to replace a single [MASK] token with multiple words. We generate the next tokens progressively by beam search and score each multitoken sequence by the harmonic mean of the probabilities of its tokens. The schematic illustration of condBERT approach is presented in Figure 4.

To evaluate the efficiency of BERT fine-tuning, we test condBERT in two scenarios:

- **zero-shot** where BERT is taken as is (with no extra fine-tuning);
- **fine-tuned** where BERT is fine-tuned on a dataset of toxic and safe sentences to acquire a style-dependent distribution, as described above.

The scenarios are different only in terms of BERT pre-training. They both use the classifier-based selection of toxic words and penalties for the toxicity of word replacements.

The strength of condBERT compared to the GPT-based method is that it does not require any parallel data. Besides that, it does not rewrite the sentence, which might be a better strategy in terms of content preservation.

5 Evaluation

To perform a comprehensive evaluation of a style transfer model, we need to make sure that it (i) changes the text style, (ii) preserves the content, and (iii) yields a grammatical sentence. The majority of works on style transfer use individual metrics to evaluate the three parameters. However, [18] points out that these three parameters are usually inversely correlated, so they need to be combined to find the balance. Our evaluation setup (individual metrics and the joint metric which combines them) follows this principle.

5.1 Style transfer accuracy

To evaluate style transfer accuracy (**STA**), we train a binary classifier $g(x_i) \rightarrow s_i$ based on RuBERT [10] that classifies text x_i into style $s_i \in \{\text{toxic}, \text{neutral}\}$. We fine-tune the RuBERT model on RuToxic dataset (see Section 6.1). It achieves the F_1 -score of 0.83 on a held-out test set. Thus, it shows a reasonable result on detection of toxic texts and can be used for evaluating the strength of style transfer. Since we want to perform the detoxification task, we expected the outputs of style transfer methods to be non-toxic. We compute the accuracy based on this assumption.

5.2 Content preservation

We approach the assessment of content preservation from two sides. First, we calculate word-based metrics: (i) the unigram word overlap (**WO**) between the tokens of the original sentence x and the style-transferred sentence y : $\frac{\text{count}(x \cap y)}{\text{count}(x \cup y)}$ and (ii) **BLEU** score, which is the ngram precision for n from 1 to 4. Secondly, we calculate the cosine similarity (**CS**) between the vector representations of the input and the output sentences. We calculate vector representations as the mean of token vector representations extracted with a fastText [3] model from RusVectors[11].³

5.3 Language quality

We use perplexity (**PPL**) to evaluate the quality of the generated sentence. As a language model for this metric, we use the ruGPT2Large⁴ model which was trained on bigger amount of content than used ruGPT3 models and was not used in our detoxGPT setups. Thus, we can claim that this model can give us the fair score for the perplexity.

5.4 Aggregated metric

Following [18], we combine the three parameters. Namely, we compute the geometric mean of STA, CS, and 1/PPL:

$$\text{GM} = (\max(\text{STA}, 0) \times \max(\text{CS}, 0) \times \max(1/\text{PPL}, 0))^{\frac{1}{3}}$$

We denote this joint metric as **GM**. Other content preservation metrics do not participate in the combination and are reported to understand the model properties better.

Although there are still discussions about the efficiency of the usage of automatic metrics for the evaluation [29] of style transfer tasks, we believe that the described metrics can adequately illustrate the strength of style transfer methods.

³<http://vectors.npl.eu/repository/20/213.zip>

⁴<https://github.com/sberbank-ai/ru-gpts#Pretraining-ruGPT2Large>

6 Experiments

We train and evaluate the two proposed models (detoxGPT and condBERT) and compare them to the baselines.

6.1 Datasets

All our methods including the *Delete* and *Retrieve* baselines require collections of toxic and non-toxic texts for training. There exist non-parallel corpora of such texts for Russian. Two corpora of toxic comments were released on Kaggle.^{5,6} We concatenate these resources and denote the joint corpus **RuToxic** dataset. It consists of 163,187 texts (31,407 (19%) toxic and 131,780 non-toxic) from the Russian social networks Odnoklassniki⁷ and Pikabu.⁸

We also use a fraction of this dataset to construct the parallel training data for detoxGPT: we select 200 toxic sentences and manually rewrite them into non-toxic ones. Besides, we use the RuToxic dataset to train toxicity weights for condBERT.

We test all models on 10,000 randomly selected toxic sentences from RuToxic. These sentences are not used for training.

6.2 Experimental Setup

For the **Delete** method, we use a manually created set of rude, obscene, and toxic words. We extend the list with word lemmas for better coverage. In **Retrieve** method we get the word vector representations from Russian *fasttext* model from the RusVectors website. The text vector representations are obtained as the mean of token vectors. We use cosine similarity as the metric of similarity between texts. For both Delete and Retrieve methods the input was preprocessed with the following steps: the input text was tokenized and obtained tokens were lemmatized with UDPipe.⁹

ruGPT3 model is available in three flavours: `small` (125m parameters with 2048 context), `medium` (350m parameters with 2048 context), and `large` (760m parameters with 2048 context). We experiment with all of them. We denote the detoxGPT models that use these ruGPT3 pretrained LMs as detoxGPT-small, detoxGPT-medium, and detoxGPT-large. ruGPT3 uses the following hyper-parameters:

- **top_k**: integer parameter that is greater or equal to 1. Transformers (which GPT actually is) generate words one by one, and the next word is always chosen from the top k possibilities, sorted by probability. We use `top_k = 3`.
- **top_p**: floating-point parameter from 0 to 1. The idea is similar to the `top_k` parameter, but the sampling is done by choosing from the smallest possible set of words whose cumulative probability exceeds the probability p . We use `top_p = 0.95`.
- **temperature** (t): floating-point parameter greater or equal to 0. It represents the degree of freedom for the model. For the higher temperatures (e.g. 100), the model can start a dialogue instead of paraphrasing, whereas for a temperature of around 1 it barely changes the sentence. We use `t = 50`.

For the few-shot and fine-tuned scenarios, we used the dataset with 200 parallel samples as described in Section 6.1.

For **condBERT** we use two setup of pre-trained weights:

- Conversational RuBERT¹⁰ from DeepPavlov [10];
- A smaller version of multilingual BERT for Russian¹¹ from Geotrend [1].

The BERT model from DeepPavlov is more commonly used for Russian language, but it is shipped without the masked LM layer that has to be trained from scratch. The BERT from Geotrend, conversely, has a pretrained LM head.

⁵<https://www.kaggle.com/blackmoon/russian-language-toxic-comments>

⁶<https://www.kaggle.com/alexandersemiletov/toxic-russian-comments>

⁷<https://ok.ru>

⁸<https://pikabu.ru>

⁹<https://ufal.mff.cuni.cz/udpipe/1/models>

¹⁰<https://huggingface.co/DeepPavlov/rubert-base-cased-conversational>

¹¹<https://huggingface.co/Geotrend/bert-base-ru-cased>

6.3 Results and Discussion

The performance of the proposed models on this data is shown in Table 1.

Method	STA \uparrow	CS \uparrow	WO \uparrow	BLEU \uparrow	PPL \downarrow	GM \uparrow
Duplicate	0.00	1.00	1.00	1.00	146.00	0.05 \pm 0.0012
Delete	0.27	0.96	0.85	0.81	263.55	0.10 \pm 0.0007
Retrieve	0.91	0.85	0.07	0.09	65.74	0.22 \pm 0.0010
detoxGPT-small						
zero-shot	0.93	0.20	0.00	0.00	159.11	0.10 \pm 0.0005
few-shot	0.17	0.70	0.05	0.06	83.38	0.11 \pm 0.0009
fine-tuned	0.51	0.70	0.05	0.05	39.48	0.20 \pm 0.0011
detoxGPT-medium						
fine-tuned	0.49	0.77	0.18	0.21	86.75	0.16 \pm 0.0009
detoxGPT-large						
fine-tuned	0.61	0.77	0.22	0.21	36.92	0.23* \pm 0.0010
condBERT						
DeepPavlov zero-shot	0.53	0.80	0.42	0.61	668.58	0.08 \pm 0.0006
DeepPavlov fine-tuned	0.52	0.86	0.51	0.53	246.68	0.12 \pm 0.0007
Geotrend zero-shot	0.62	0.85	0.54	0.64	237.46	0.13 \pm 0.0009
Geotrend fine-tuned	0.66	0.86	0.54	0.64	209.95	0.14 \pm 0.0009

Table 1: The results of evaluation of proposed detoxification approaches. **STA**: Style transfer accuracy. **CS**: Cosine similarity. **WO**: Word overlap rate. **PPL**: Perplexity. **GM**: Geometric mean. The larger \uparrow (or lower \downarrow), the better. Gray numbers show that a method significantly fails to preserve the content. The values **in bold** are the best scores. The asterisk * denotes the improvement over the **Retrieve** baseline that is statistically significant at $p \leq 0.01$. The standard deviations of **GM** are calculated by bootstrapping the test dataset.

The baseline approaches represent the two extremes: while **Delete** gains a low STA and high content similarity, the **Retrieve** method, on the contrary, achieves a relatively high STA with extremely low WO and BLEU. These results are natural since the Delete method only eliminates toxic words and leaves the rest of the sentence intact, which results in high word-based similarity. At the same time, such deletion of words often ruins the sentence structure and results in high PPL. The Retrieve method always outputs only non-toxic, fully human-readable sentences; this strategy achieves a high STA score and the highest GM score between baselines. However, the content of such sentences is unpredictable and usually very different from the original input.

We experiment with *zero-shot*, *few-shot*, and *fine-tuned* setups for the three **detoxGPT** model versions as described in Section 4.2. However, the quality of the output of the *zero-shot* and *few-shot* scenarios is poor for all models. Thus, we report the results of *zero-shot*, *few-shot* only for the detoxGPT-small model to illustrate the difference in scores. Table 1 shows that content similarity and fluency of both *zero-shot* and *few-shot* models are lower than those of the baselines. The *zero-shot* method manages to reach high style accuracy by generating completely irrelevant texts which happen to be mostly non-toxic. As a result, we do not take into account its results in comparison with other approaches. On the other hand, when fine-tuned on only 200 samples, detoxGPT models outperform the baselines. The best results are achieved by the **detoxGPT-large** model. It reaches the highest values for all metrics (and the lowest for PPL which stands for the highest naturalness) including the joint GM score.

The **condBERT**-based models also outperform the **Delete** baseline, but fall short of the **Retrieve** baseline due to lower fluency. The condBERT models based on Geotrend pre-trained BERT model show better performance than DeepPavlov setup in general. The reason is the pre-trained language model part in Geotrend RuBERT. For DeepPavlov setup these weights of the model were not pre-trained and were initialized randomly. The comparison with detoxGPT is ambiguous: condBERT attains the highest STA score and larger values of content similarity metrics, while the quality of language is quite low. The high

style transfer accuracy is explained by the fact that the replacement words in condBERT are explicitly penalized for toxicity. Nevertheless, the model can generate similar toxic or rude words, making its style transfer quality imperfect. At the same time, as detoxGPT generates the sentences from scratch, we cannot control the content preservation (the models occasionally change it entirely). In contrast, the condBERT model changes only a small fraction of words. However, the downside of condBERT is the loss of naturalness. GPT-based models are free to choose every single word in the output, usually yielding a natural sentence. Conversely, when replacing only single words (as condBERT does), it is difficult to control the overall grammar and fluency. Furthermore, condBERT, for instance, in DeepPavlov setup sometimes substitutes words with special symbols such as [UNK].

Model	Sentence
Input	не дай бог моя дочь так оденется убью нахуй палкой (If, God forbid, my daughter goes out dressed like this, I'll fucking kill her with a stick)
Delete	не дай бог моя дочь так оденется убью палкой (If, God forbid, my daughter goes out dressed like this, I'll kill her with a stick)
Retrieve	не бросайте угла родного одной мы лежали больнице палате в в в те девчонкой была молодой годы (don't abandon your native corner same hospital we were ward in in in those girl was young years)
detoxGPT-small	<u>подскажите, у меня дочь так оденется палкой.</u> (Could you give me an advise, my daughter gets dressed as a stick.)
detoxGPT-medium	да не торопись вы так, бабушка. (Granny, don't rush, take your time.)
detoxGPT-large	<u>извините, но мне вас очень жаль, если вы начнете так одеваться.</u> (Excuse me, but I feel sorry for you if you start getting dressed up like this.)
condBERT	не дай бог моя дочь так оденется станки палкой (God forbid, my daughter gets dressed up like this machine stick.)
Input	что ты читала из его произведений? дура.а если не читала,то прочти. (Which of his works have you read? You are an idiot. If you haven't read any, go read.)
Delete	что ты читала из его произведений дура.а если не читала,то прочти (Which of his works have you read you idiot. If you haven't read any, go read.)
Retrieve	одно не пойму, вы все реально думаете, что они читают что то в ок или в вк? (What I can't understand - do you all really think they read anything in Odnoklassniki or Vkontakte?)
detoxGPT-small	<u>вы читали его произведения? не читали.</u> (Have you read any of his works? No, you haven't)
detoxGPT-medium	его посадить, извините, но мне вас очень жаль, если для вас оскорбления - норма. (Putting him in jail, excuse me, but I feel sorry for you if insult is normal for you.)
detoxGPT-large	<u>вы читали его произведения? недалекие. попробуйте почитать посты его друзей</u> (Have you read any of his works? You are narrow-minded. Try reading his friends' posts.)
condBERT	<u>что ты читала из его произведений ? пусть . а если не читала , то прочти .</u> (Which of his works have you read? Okay. If you haven't read any, go read.)
Input	тебя вношу в список долбаебов. (I'm adding you to the list of buttheads.)
Delete	тебя вношу в список (I'm adding you to a list)
Retrieve	надо обнародовать этот список. чего молчать. (We should make this list public. Why keeping silence?)
detoxGPT-small	<u>Вас вношу в список людей, нуждающихся в улучшении.</u> (I'm adding you to the list of people who need being improved.)
detoxGPT-medium	<u>Вас вношу в список людей, которые вносят вклад в мой тред.</u> (I'm adding you to the list of people who contribute to this thread.)
detoxGPT-large	<u>Вас вношу в список людей, которые не соответствуют вашим ожиданиям.</u> (I'm adding you to the list of people who don't meet your expectations.)
condBERT	тебя вношу в список до (I'm adding you to the list of to.)

Table 2: Examples of Russian texts detoxification by proposed approaches. For detoxGPT models, the results of fine-tuned setup are presented. For condBERT model, the results of Geotrend fine-tuned setup are presented. The rude words used in sentences have no goal to abuse the reader, they are just an illustration of real-life toxic texts. The best outputs for each example according to a human judgment are underlined.

Table 2 shows the example outputs of the proposed models and the baselines. All the examples by detoxGPT and condBERT models were generated via the *fine-tuned* scenario. The examples demonstrate the trends described above: condBERT sometimes makes an inappropriate replacement, and detoxGPT tends to output sentences not related to the input. Nevertheless, in most cases, at least one of the detoxGPT models provides a sensible answer. Interestingly, although detoxGPT-large performs best according to the metrics, the manual analysis shows that its superiority is not always evident.

7 Conclusion

We presented the first study of text detoxification for the Russian language. We conducted experiments with detoxification methods based on different principles: (i) detoxGPT model is trained on a parallel corpus and rewrites the sentence, and (ii) condBERT is trained on non-parallel data and replaces individual toxic words with non-toxic synonyms. We described the evaluation setup, which includes the training and test data and the evaluation metrics. We evaluated the proposed methods and compare them to three simple baselines.

The best aggregated score is achieved by detoxGPT. While condBERT shows the highest style transfer accuracy, it performs worse in naturalness preservation. However, for both methods, there is room for improvement. The detoxGPT-based models could benefit from a larger parallel corpus and more careful tuning of hyperparameters, while for condBERT, more advanced word selection strategies can increase the quality.

As a result, there is no single method that outperforms others according to all parameters of the evaluation. Sometimes it is enough to delete obscene words from the text, whereas in other cases, they should be replaced with their non-toxic synonyms. Finally, some texts can be detoxified only if fully reformulated. Thus, the most promising direction of future work would be to combine all presented strategies and apply them based on the nature of toxicity in particular sentences.

We provide all code and data used for training and evaluation online.¹²

Acknowledgements

This work was conducted under the framework of the joint Skoltech-MTS laboratory. We are grateful to the anonymous reviewers for their helpful suggestions. Besides, we thank Alexey Shevtsov and Alexander Nevarko who conducted the first version of experiments with ruGPT as a part of their Deep Learning course final project at Skoltech.

References

- [1] Amine Abdaoui, Camille Pradel, and Grégoire Sigel. Load what you need: Smaller versions of multilingual BERT. In *Proceedings of SustaiNLP: Workshop on Simple and Efficient Natural Language Processing*, pages 119–123, Online, November 2020. Association for Computational Linguistics.
- [2] Nikolay Arefyev, Boris Sheludko, Alexander Podolskiy, and Alexander Panchenko. Always keep your target in mind: Studying semantics and improving performance of neural lexical substitution. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1242–1255, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics.
- [3] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *Trans. Assoc. Comput. Linguistics*, 5:135–146, 2017.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina 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)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.

¹²<https://github.com/skoltech-nlp/rudetoxifier>

- [5] Ashwin Geet D'Sa, Irina Illina, and Dominique Fohr. Towards non-toxic landscapes: Automatic toxic comment detection using DNN. In *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*, pages 21–25, Marseille, France, May 2020. European Language Resources Association (ELRA).
- [6] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2414–2423. IEEE Computer Society, 2016.
- [7] Di Jin, Zhijing Jin, Zhiting Hu, Olga Vechtomova, and Rada Mihalcea. Deep learning for text style transfer: A survey. *CoRR*, abs/2011.00416, 2020.
- [8] Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. Disentangled representation learning for non-parallel text style transfer. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 424–434, Florence, Italy, July 2019. Association for Computational Linguistics.
- [9] Kalpesh Krishna, John Wieting, and Mohit Iyyer. Reformulating unsupervised style transfer as paraphrase generation. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu, editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 737–762. Association for Computational Linguistics, 2020.
- [10] Yuri Kuratov and Mikhail Arkhipov. Adaptation of deep bidirectional multilingual transformers for russian language. *CoRR*, abs/1905.07213, 2019.
- [11] Andrey Kutuzov and Elizaveta Kuzmenko. *WebVectors: A Toolkit for Building Web Interfaces for Vector Semantic Models*, pages 155–161. Springer International Publishing, Cham, 2017.
- [12] Juncen Li, Robin Jia, He He, and Percy Liang. Delete, retrieve, generate: a simple approach to sentiment and style transfer. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1865–1874, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.
- [13] Aman Madaan, Amrith Setlur, Tanmay Parekh, Barnabás Póczos, Graham Neubig, Yiming Yang, Ruslan Salakhutdinov, Alan W. Black, and Shrimai Prabhunoye. Politeness transfer: A tag and generate approach. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1869–1881. Association for Computational Linguistics, 2020.
- [14] David D. McDonald and James Pustejovsky. A computational theory of prose style for natural language generation. In Maghi King, editor, *EACL 1985, 2nd Conference of the European Chapter of the Association for Computational Linguistics, March 27-29, 1985, University of Geneva, Geneva, Switzerland*, pages 187–193. The Association for Computer Linguistics, 1985.
- [15] Igor Melnyk, Cícero Nogueira dos Santos, Kahini Wadhawan, Inkit Padhi, and Abhishek Kumar. Improved neural text attribute transfer with non-parallel data. *CoRR*, abs/1711.09395, 2017.
- [16] Cicero Nogueira dos Santos, Igor Melnyk, and Inkit Padhi. Fighting offensive language on social media with unsupervised text style transfer. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 189–194, Melbourne, Australia, July 2018. Association for Computational Linguistics.
- [17] Endang Wahyu Pamungkas and Viviana Patti. Cross-domain and cross-lingual abusive language detection: A hybrid approach with deep learning and a multilingual lexicon. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 363–370, Florence, Italy, July 2019. Association for Computational Linguistics.
- [18] Richard Yuanzhe Pang and Kevin Gimpel. Unsupervised evaluation metrics and learning criteria for non-parallel textual transfer. In Alexandra Birch, Andrew M. Finch, Hiroaki Hayashi, Ioannis Konstas, Thang Luong, Graham Neubig, Yusuke Oda, and Katsuhito Sudoh, editors, *Proceedings*

- of the 3rd Workshop on Neural Generation and Translation@EMNLP-IJCNLP 2019, Hong Kong, November 4, 2019, pages 138–147. Association for Computational Linguistics, 2019.
- [19] Reid Pryzant, Richard Diehl Martinez, Nathan Dass, Sadao Kurohashi, Dan Jurafsky, and Diyi Yang. Automatically neutralizing subjective bias in text. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 480–489. AAAI Press, 2020.
- [20] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [21] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020.
- [22] Sudha Rao and Joel Tetreault. Dear sir or madam, may I introduce the GYAFC dataset: Corpus, benchmarks and metrics for formality style transfer. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 129–140, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.
- [23] Anna Schmidt and Michael Wiegand. A survey on hate speech detection using natural language processing. In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, pages 1–10, Valencia, Spain, April 2017. Association for Computational Linguistics.
- [24] Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi S. Jaakkola. Style transfer from non-parallel text by cross-alignment. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6830–6841, 2017.
- [25] Alexey Tikhonov and Ivan P. Yamshchikov. What is wrong with style transfer for texts? *CoRR*, abs/1808.04365, 2018.
- [26] Minh Tran, Yipeng Zhang, and Mohammad Soleymani. Towards a friendly online community: An unsupervised style transfer framework for profanity redaction. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2107–2114, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics.
- [27] Xing Wu, Shangwen Lv, Liangjun Zang, Jizhong Han, and Songlin Hu. Conditional BERT contextual augmentation. In João M. F. Rodrigues, Pedro J. S. Cardoso, Jânio M. Monteiro, Roberto Lam, Valeria V. Krzhizhanovskaya, Michael Harold Lees, Jack J. Dongarra, and Peter M. A. Sloot, editors, *Computational Science - ICCS 2019 - 19th International Conference, Faro, Portugal, June 12-14, 2019, Proceedings, Part IV*, volume 11539 of *Lecture Notes in Computer Science*, pages 84–95. Springer, 2019.
- [28] Xing Wu, Tao Zhang, Liangjun Zang, Jizhong Han, and Songlin Hu. "mask and infill" : Applying masked language model to sentiment transfer. *CoRR*, abs/1908.08039, 2019.
- [29] Ivan P. Yamshchikov, Viacheslav Shibaev, Nikolay Khlebnikov, and Alexey Tikhonov. Style-transfer and paraphrase: Looking for a sensible semantic similarity metric. *CoRR*, abs/2004.05001, 2020.