

SmurfCat at PAN 2024 TextDetox: Alignment of Multilingual Transformers for Text Detoxification

Notebook for the PAN Lab at CLEF 2024

Elisei Rykov^{1,*}, Konstantin Zaytsev^{2,*}, Ivan Anisimov¹ and Alexandr Voronin¹

¹Skolkovo Institute of Science and Technology, Russia

²HSE University, Russia

Abstract

This paper presents a solution for the *Multilingual Text Detoxification* task in the PAN-2024 competition of the *SmurfCat* team. Using data augmentation through machine translation and a special filtering procedure, we collected an additional multilingual parallel dataset for text detoxification. Using the obtained data, we fine-tuned several multilingual sequence-to-sequence models, such as mT0 and Aya, on a text detoxification task. We applied the ORPO alignment technique to the final model. Our final model has only 3.7 billion parameters and achieves state-of-the-art results for the Ukrainian language and near state-of-the-art results for other languages. In the competition, our team achieved first place in the automated evaluation with a score of 0.52 and second place in the final human evaluation with a score of 0.74.

Keywords

PAN 2024, Multilingual Text Detoxification, mT0, ORPO

1. Introduction

Multilingual text detoxification is a challenging subtask within text style transfer. The most difficult part is the adaptation of such a system to low-resource languages. The concept of PAN-2024 Multilingual Text Detoxification Task [1, 2] is to develop a multilingual text detoxification system for 9 languages: Amharic, Arabic, German, Spanish, Hindi, Chinese, Russian, Ukrainian and English.

This paper describes the solution of the *SmurfCat* team, which achieved first place with an average score of 0.52 in the automatic evaluation and second place with a score of 0.74 in the manual human evaluation. Our solution is based on the mT0 model family [3], which has powerful multilingual capabilities. We fine-tuned all our selected models to each language of the competition, and applied various data augmentation techniques. To improve detoxification, we performed hypothesis filtering using the diverse beam search algorithm [4]. Finally, we applied ORPO [5] alignment to enforce model predictions. Our 3.7-billion-parameter language model demonstrates state-of-the-art results for Ukrainian and near state-of-the-art results for other languages. We published the final best-performing model on the HuggingFace Hub¹. You can also find the training scripts and the extended data on GitHub².

The rest of the paper is organized as follows: Section 2 discusses data augmentation strategies, Section 3 describes our final solution, and Section 4 presents the results and discussion.

CLEF 2024: Conference and Labs of the Evaluation Forum, September 09–12, 2024, Grenoble, France

*These authors contributed equally.

✉ Elisei.Rykov@skoltech.ru (E. Rykov); kzaytsev@hse.ru (K. Zaytsev); Ivan.Anisimov@skoltech.ru (I. Anisimov); Alexandr.Voronin@skoltech.ru (A. Voronin)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://hf.co/s-nlp/mt0-xl-detox-orpo>

²<https://github.com/s-nlp/multilingual-transformer-detoxification>

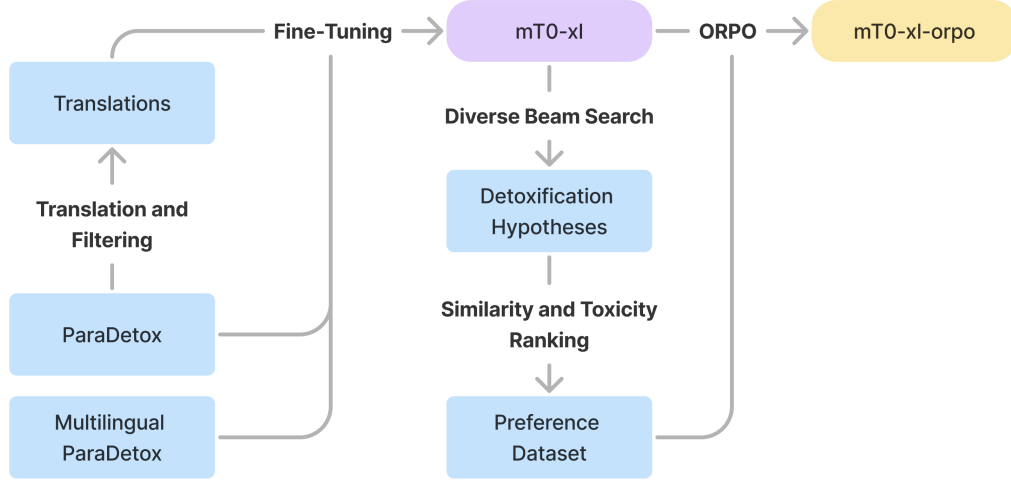


Figure 1: An overview of our approach. We used different datasets, fine-tuned the whole mT0-XL model and finally performed the ORPO alignment step.

2. Data

Initially, there were not many parallel datasets for the multilingual detoxification task. More precisely, primarily only the Russian³ and English⁴ ParaDetox datasets were available, with 11 100 and 19 700 samples respectively. During the competition, the organizers published a small human-annotated Multilingual ParaDetox⁵ for all languages, containing only 400 samples per language.

Nevertheless, we decided to augment the provided data by automatic translation from English to other languages. To translate the original English data, we used a GoogleTranslator model from deep_translator⁶ Python package. We chose API over some of the more advanced machine translation models because of its speed and simplicity. Also, there are not as many translators for low-resource languages like Amharic. As a result, we obtained an additional 19 700 samples for each language.

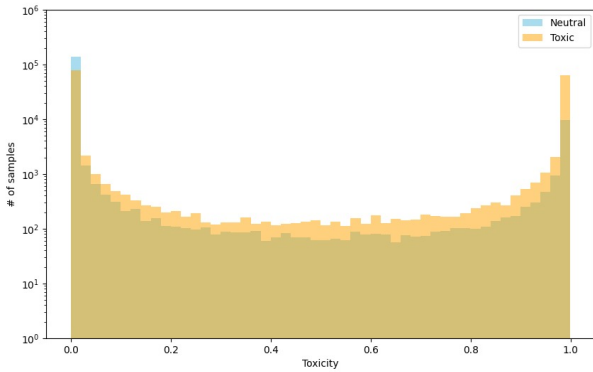


Figure 2: Toxicity of translations

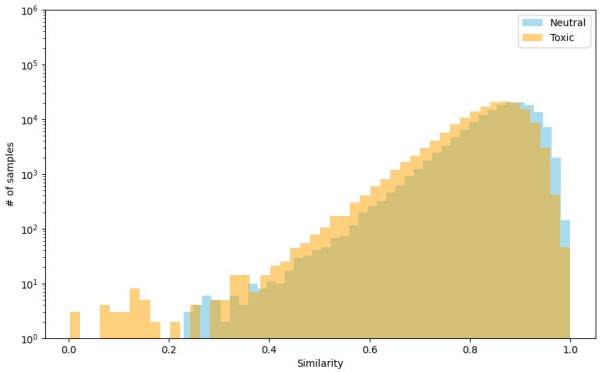


Figure 3: Similarity of translations

Since translation is often imperfect, we decided to perform a specific post-processing procedure. In general, we checked the preservation of meaning after translation and the toxicity of the translated data. First, we used the LaBSE [6] model to evaluate the similarity between translated pairs. Second, we applied XLM-R⁷ toxicity classifier to check whether toxic sentences were still toxic after translation

³https://huggingface.co/datasets/s-nlp/ru_paradeto

⁴<https://huggingface.co/datasets/s-nlp/paradeto>

⁵https://huggingface.co/datasets/textdetox/multilingual_paradeto

⁶<https://pypi.org/project/deep-translator/>

⁷<https://huggingface.co/textdetox/xlmr-large-toxicity-classifier>

and vice versa.

A distribution for both of two measures is shown on Figures 2, 3. For most samples, similarity between original and translated samples was high enough that many samples preserved their meaning. Regarding toxicity, many neutral sentences became toxic after translation, and many toxic sentences became neutral. For toxicity, we set a threshold parameter to 0.9 for toxic sentences and 0.1 for neutral sentences. The similarity threshold was set to 0.8 for all sentences.

After all filtering steps, 40 500 pairs of neutral and toxic sentences were obtained. A more precise statistic of how many samples remain after filtering is given in Table 1. According to the statistics, Amharic lost the most samples during filtering.

Table 1

Statistics of number of remaining samples after filtering.

Language	Amharic	Arabic	German	Spanish	Hindi	Russian	Ukrainian	Chinese
# of samples	1 323	3 190	7 511	7 555	4 844	7 458	5 350	3 274

Our final dataset mixture is shown in Table 2. In total, 74 900 samples were used in the training process.

3. Method

In this section, we describe our prior method, related to fine-tuning and optimization of Language Models on the text detoxification task.

3.1. Supervised fine-tuning

As a main approach, we choose fine-tuning of various multilingual LMs. As we suggest, the most promising models for the further fine-tuning were LMs from mT0 family. It is a family of sequence-to-sequence Transformer models initialized from mT5 [7]. We considered that sequence-to-sequence modeling would be more preferable in case of the text detoxification task. The mT0 family was chosen because of its strong multilingual capabilities, so these models were adapted to each language of the competition. We also experimented with the novel Aya-101 model [8]: a fine-tuned mT5-x1 model on a multilingual instructions.

All models were tuned in an almost similar way. The learning rate was set to $1e-5$, the global batch size to 8, and the weight decay to 0.01. The cosine scheduler was used for training. In total, 4 all models were trained during 4 epochs. All other training parameters were default according to HuggingFace Seq2SeqTrainer. The only difference is that for the mT0-XL we updated the weights of the whole model because our computing resources allowed it. In case of a larger model like Aya-101 or mT0-XXL, only the LoRA adapter was trained. The setup of the LoRA adapter was as follows: r and lora alpha were set to 32, lora dropout parameter to 0.1, other parameters were default. The best model was selected according to the validation loss.

To enforce the in-context abilities of the models, we added a specific prefix to each toxic sentence depending on the language. As a result, we passed toxic sentences with special prefix prompt into the model during training.

3.2. The Best Candidate Choice

During inference, we generated 10 hypotheses and selected 5 most likely ones using diverse beam search. The number of beams was set to 10 with 5 beam groups, the diversity penalty was 2.5, the repetition penalty was 1.2. To select the best choice, we calculated a relevance metric using a product of similarity and toxicity scores. Similarity was calculated using LaBSE embeddings, and toxicity was measured using the xlm-roberta-large toxicity classifier. As relevance scores were calculated, we selected then the best candidate according to the highest score.

Table 2

Training dataset mixture.

Dataset	# of samples
En-ParaDetox	19 700
Ru-ParaDetox	11 100
Translations	40 500
Multilingual ParaDetox	3 600
Total	74 900

3.3. ORPO

Once the models were fine-tuned, we decided to further tune the model for best performance using Odds Ratio Preference Optimization (ORPO) approach. This optimization does not need a reference model like it is in DPO [9]. Alignment was employed on the unseen test dataset.

As a preference dataset, we generated hypotheses using diverse beam search on the samples from the test set and annotated them using the relevance score described above. Only candidates with the highest relevance scores were selected as the chosen ones, and all others were selected as the rejected samples.

The final ORPO data set for alignment contained the prompt (toxic sentence), the rejected sample (negative candidate), and the selected sample (best candidate). Table 3 shows a small sample of the dataset. Since the dataset was collected, we trained the model on the dataset using the same parameters used to train the other models. Since ORPO uses the beta parameter, it was set to 0.1. For the final submission, we used an aligned model with the algorithm described above to select the best candidate.

Table 3

Samples for the ORPO tuning.

Prompt	Chosen	Rejected
Detoxify: She really can't be such a bitch, can she???	She really can't be such a bad person, can she???	really can't be such a bad person, can she???
Detoxify: hopefully turks are not dumb enough to hand it to him	Turkish people are not naive enough to hand it to him.	hopefully turks are not naive enough to hand it to him.
Detoxify: again , give me the name of the store or fuck off , liar .	again, give me the name of the store or go away	again, give me the name of the store
Detoxify: Nat is just a piece of shit, ignore him.	Nate is just not good, ignore him.	Nat is just a bad person, ignore him

4. Results

The final results of the automatic evaluation are shown in the Table 4. The mT0-XL with ORPO alignment showed the best performance among all approaches from the leaderboard for all languages. Compared to mT0-XL, a model before ORPO alignment, ORPO slightly improved the performance of the model, increasing the average results by 0.01 points. Surprisingly, the larger models are not the best. For example, the mT0-XXL model with 13B parameters performed even worse than the mT0-XL model with only 3.7B parameters. Aya-101, an mT5-XXL model additionally tuned to instructional data for different languages, performed worse than other models. Since Aya-101 and mT0-XXL performed even worse on mT0-XL, we did not perform an ORPO alignment step for these models. Considering other teams on the automatic evaluation, our checkpoints, mainly mT0-XL-ORPO and mT0-XL, are the two best performing approaches for all languages except the Chinese language.

Table 4

The results of the automatic evaluation. The teams with the best scores were selected for the table. *Joint* is given as the evaluation metric.

Team	Language									Avg J
	Amharic	Arabic	German	English	Spanish	Hindi	Russian	Ukrainian	Chinese	
<i>Our (mT0-XL-ORPO)</i>	0.378	0.626	0.678	0.602	0.562	0.355	0.634	0.692	0.178	0.523
<i>Our (mT0-XL)</i>	<u>0.374</u>	<u>0.617</u>	<u>0.669</u>	<u>0.593</u>	<u>0.555</u>	<u>0.352</u>	<u>0.628</u>	<u>0.686</u>	<u>0.165</u>	<u>0.515</u>
<i>Our (mT0-XXL-LoRA)</i>	<u>0.361</u>	<u>0.594</u>	<u>0.639</u>	<u>0.591</u>	<u>0.548</u>	<u>0.345</u>	<u>0.605</u>	<u>0.660</u>	<u>0.159</u>	<u>0.500</u>
nikita.sushko	0.328	0.575	0.592	0.553	0.480	0.241	0.570	0.668	<u>0.176</u>	0.465
VitalyProtasov	0.311	0.523	0.502	0.531	0.472	0.320	0.542	0.629	0.175	0.445
erehulka	0.287	0.536	0.575	0.543	0.497	0.185	0.529	0.602	0.160	0.435
<i>Our (Aya-101-LoRA)</i>	<u>0.301</u>	<u>0.526</u>	<u>0.530</u>	<u>0.529</u>	<u>0.475</u>	<u>0.223</u>	<u>0.541</u>	<u>0.586</u>	<u>0.108</u>	<u>0.424</u>
ansafonov	0.270	0.456	0.362	0.506	0.319	0.133	0.507	0.328	0.178	0.340

The Table 5 shows human evaluation results. Our detoxification model for Ukrainian achieved the highest human evaluation score by a wide margin, indicating that our approach is the state-of-the-art for this language. Overall, our best performing checkpoint is the top-2 approach according to the human evaluation by the averaged *Joint* metric.

Table 5

The results of the human evaluation. The teams with the best scores were selected for the table. *Joint* is given as the evaluation metric.

Team	Language									Avg J
	Amharic	Arabic	German	English	Spanish	Hindi	Russian	Ukrainian	Chinese	
Human Reference	0.85	0.82	0.71	0.88	0.79	0.97	0.80	0.90	0.93	0.85
SomethingAwful	<u>0.71</u>	0.74	0.89	0.86	0.83	<u>0.86</u>	0.84	0.69	0.53	0.77
<i>Our (mT0-XL-ORPO)</i>	<u>0.71</u>	<u>0.82</u>	<u>0.70</u>	<u>0.83</u>	<u>0.73</u>	<u>0.68</u>	<u>0.76</u>	0.84	<u>0.60</u>	<u>0.74</u>
VitalyProtasov	0.68	0.79	0.77	0.69	<u>0.81</u>	0.87	0.73	0.67	0.49	0.72
nikita.sushko	0.68	0.89	0.79	0.70	0.62	0.84	0.74	0.67	0.47	0.71
erehulka	0.69	0.78	<u>0.85</u>	0.88	0.71	0.52	0.65	0.63	<u>0.68</u>	0.69
mkrisnai	0.49	0.63	0.70	<u>0.89</u>	0.83	0.73	<u>0.78</u>	<u>0.73</u>	0.34	0.68
d1n910	0.61	0.44	0.77	0.91	0.77	0.34	0.71	0.50	0.84	0.65
ZhongyuLuo	0.72	0.49	0.01	0.73	0.52	0.49	0.68	0.42	0.56	0.51

5. Conclusion

In conclusion, our system demonstrated a strong pipeline for augmenting training data for low-resource languages and further fine-tuning a relatively small 3.7 billion parameter language model for the text detoxification task. Our future research may consider how to adapt text detoxification capabilities from high-resource languages to low-resource languages without translation, as machine translation for low-resource languages often shows low quality. A further direction for investigation is the interpretability of models, specifically the understanding of which tokens have been replaced by the model through the text detoxification process and the rationale behind this.

References

- [1] J. Bevendorff, X. B. Casals, B. Chulvi, D. Dementieva, A. Elnagar, D. Freitag, M. Fröbe, D. Korenčić, M. Mayerl, A. Mukherjee, A. Panchenko, M. Potthast, F. Rangel, P. Rosso, A. Smirnova, E. Stamatatos, B. Stein, M. Taulé, D. Ustalov, M. Wiegmann, E. Zangerle, Overview of PAN 2024: Multi-Author Writing Style Analysis, Multilingual Text Detoxification, Oppositional Thinking Analysis, and Generative AI Authorship Verification, in: Experimental IR Meets Multilinguality, Multimodality,

- and Interaction. Proceedings of the Fourteenth International Conference of the CLEF Association (CLEF 2024), Lecture Notes in Computer Science, Springer, Berlin Heidelberg New York, 2024.
- [2] D. Dementieva, D. Moskovskiy, N. Babakov, A. A. Ayele, N. Rizwan, F. Schneider, X. Wang, S. M. Yimam, D. Ustalov, E. Stakovskii, A. Smirnova, A. Elnagar, A. Mukherjee, A. Panchenko, Overview of the multilingual text detoxification task at pan 2024, in: G. Faggioli, N. Ferro, P. Galuščáková, A. G. S. de Herrera (Eds.), Working Notes of CLEF 2024 - Conference and Labs of the Evaluation Forum, CEUR-WS.org, 2024.
 - [3] N. Muennighoff, T. Wang, L. Sutawika, A. Roberts, S. Biderman, T. Le Scao, M. S. Bari, S. Shen, Z. X. Yong, H. Schoelkopf, X. Tang, D. Radev, A. F. Aji, K. Almubarak, S. Albanie, Z. Alyafeai, A. Webson, E. Raff, C. Raffel, Crosslingual generalization through multitask finetuning, in: A. Rogers, J. Boyd-Graber, N. Okazaki (Eds.), Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Toronto, Canada, 2023, pp. 15991–16111. URL: <https://aclanthology.org/2023.acl-long.891>. doi:10.18653/v1/2023.acl-long.891.
 - [4] A. K. Vijayakumar, M. Cogswell, R. R. Selvaraju, Q. Sun, S. Lee, D. Crandall, D. Batra, Diverse beam search: Decoding diverse solutions from neural sequence models, 2017. URL: <https://openreview.net/forum?id=HJV1zP5xg>.
 - [5] J. Hong, N. Lee, J. Thorne, Orpo: Monolithic preference optimization without reference model, 2024. URL: <https://arxiv.org/abs/2403.07691>. arXiv:2403.07691.
 - [6] F. Feng, Y. Yang, D. Cer, N. Arivazhagan, W. Wang, Language-agnostic BERT sentence embedding, in: S. Muresan, P. Nakov, A. Villavicencio (Eds.), Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022, Association for Computational Linguistics, 2022, pp. 878–891. URL: <https://doi.org/10.18653/v1/2022.acl-long.62>. doi:10.18653/v1/2022.ACL-LONG.62.
 - [7] L. Xue, N. Constant, A. Roberts, M. Kale, R. Al-Rfou, A. Siddhant, A. Barua, C. Raffel, mt5: A massively multilingual pre-trained text-to-text transformer, in: North American Chapter of the Association for Computational Linguistics, 2020. URL: <https://api.semanticscholar.org/CorpusID:225040574>.
 - [8] A. Üstün, V. Aryabumi, Z.-X. Yong, W.-Y. Ko, D. D’souza, G. Onilude, N. Bhandari, S. Singh, H.-L. Ooi, A. Kayid, F. Vargus, P. Blunsom, S. Longpre, N. Muennighoff, M. Fadaee, J. Kreutzer, S. Hooker, Aya model: An instruction finetuned open-access multilingual language model, arXiv preprint arXiv:2402.07827 (2024).
 - [9] R. Rafailov, A. Sharma, E. Mitchell, C. D. Manning, S. Ermon, C. Finn, Direct preference optimization: Your language model is secretly a reward model, in: Thirty-seventh Conference on Neural Information Processing Systems, 2023. URL: <https://openreview.net/forum?id=HPuSIXJaa9>.