

Improving Moderation of Online Discussions via Interpretable Neural Models

Andrej Švec¹, Matúš Pikuliak², Marián Šimko², Mária Bieliková²

Slovak University of Technology in Bratislava, Bratislava, Slovakia

Faculty of Informatics and Information Technologies

¹andy.swec@gmail.com

²{matus.pikuliak, marian.simko, maria.bielikova}@stuba.sk

Abstract

Growing amount of comments make online discussions difficult to moderate by human moderators only. Antisocial behavior is a common occurrence that often discourages other users from participating in discussion. We propose a neural network based method that partially automates the moderation process. It consists of two steps. First, we detect inappropriate comments for moderators to see. Second, we highlight inappropriate parts within these comments to make the moderation faster. We evaluated our method on data from a major Slovak news discussion platform.

1 Introduction

Keeping the discussion on a website civil is important for user satisfaction as well as for legal reasons (European court of human rights, 2015). Manually moderating all the comments might be too time consuming. Larger news and discussion websites receive hundreds of comments per minute which might require huge moderator teams. In addition it is easy to overlook inappropriate comments due to human error. Automated solutions are being developed to reduce moderation time requirements and to mitigate the error rate.

In this work we propose a neural network based method to speed up the moderation process. First, we use trained classifier to automatically detect inappropriate comments. Second, a subset of words is selected with a method by (Lei et al., 2016) based on reinforcement learning. These selected words should form a rationale why a comment was classified as inappropriate by our model. Selected words are then highlighted for moderators so they can quickly focus on problematic parts of comments. We also managed to evaluate our solution on a major dataset (millions of comments) and in real world conditions at an important Slovak news discussion platform.

2 Related work

Inappropriate comments detection. There are various approaches to detection of inappropriate comments in online discussions (Schmidt and Wiegand, 2017). The most common approach is to detect inappropriate texts through machine learning. Features used include bag of words (Burnap and Williams, 2016), lexicons (Gitari et al., 2015), linguistic, syntactic and sentiment features (Nobata et al., 2016), Latent Dirichlet Allocation features (Zhong et al., 2016) or comment embeddings (Djuric et al., 2015). Deep learning was also considered to tackle this issue (Badjatiya et al., 2017; Mehdad and Tetreault, 2016).

Apart from detecting inappropriate texts, multiple works focus on detecting users that should be banned (Cheng et al., 2015) by analyzing their posts and their activity in general (Adler et al., 2011; Ribeiro et al., 2018a), their relationships with other users (Ribeiro et al., 2018b) and the reaction of other users (Cheng et al., 2014) or moderators (Cheng et al., 2015) towards them.

Interpreting neural models. Interpretability of machine learning models is common requirement when deploying the models to production. In our case moderators would like to know why was the comment marked as inappropriate.

Most of the works deal with interpretability of computer vision models (Zeiler and Fergus, 2014), but progress in interpretable text processing was also made. Several works try to analyze the dynamics of what is happening inside the neural network. Karpathy et al. (2015); Aubakirova and Bansal (2016) focus on memory cells activations. Li et al. (2016a) compute how much individual input units contribute to the final decision. Other techniques rely on attention mechanisms (Yang et al., 2016), contextual decomposition (Murdoch et al., 2018), representa-

tion erasure (Li et al., 2016b) or relevance propagation (Arras et al., 2017). Our work uses the method by Lei et al. (2016), which selects coherent subset of words responsible for neural network decision. The authors use this model to explain multi-aspect sentiment analysis over beer reviews and information retrieval over CQA system.

In the domain of detecting antisocial behavior Pavlopoulos et al. (2017) used attention mechanism to interpret existing model. Their work is the most relevant to ours, but we use other datasets as well as other, more explicit, technique for model interpretation.

3 Interpretable Neural Moderation

We propose a method to speed up the moderation process of online discussions. It consists of two steps:

1. We detect inappropriate comments. This is a binary classification problem. Comments are sorted by model confidence and shown to the moderators. In effect after this initial filtering moderators work mostly with inappropriate comments which improves their efficiency.
2. We highlight critical parts of inappropriate comments to convince moderators that selected comments are indeed harmful. The moderators can then focus on these highlighted parts instead of reading the whole comment.

3.1 Step 1: Inappropriate comments detection

We approach inappropriate comments detection as a binary classification problem. Each comment is either appropriate or inappropriate. We use recurrent neural network that takes sequence of word embeddings as input. The final output is then used to predict the probability of comment being inappropriate. We use RCNN recurrent cells (Bazilay et al., 2016) instead of more commonly used LSTM cells as they proved to be faster to train with practically identical results. This part of our method is trained in supervised fashion using Adam optimization algorithm.

3.2 Step 2: Inappropriate parts highlighting

Our method implements (Lei et al., 2016). It can learn to select the words responsible for a decision of a neural network called *rationale* without the need for word level annotations in the data.

The model processes comment word embeddings x and generates two outputs: binary flags z representing selection of individual words into rationale which is marked (z, x) and y being probability distribution over classes appropriate / inappropriate. The model is composed of two modules: generator *gen* and classifier *clas* called also encoder in the original work.

Generator *gen*. The role of generator is to select words that are responsible for a comment being in/appropriate. On its output layer it generates probabilities of selection for each word $p(z|x)$. A well trained model assigns high probability scores to words that should form the rationale and low scores to the rest. In the final step these probabilities are used to sample binary selections z . The sampling layer is called Z-layer.

Due to sampling in Z-layer *gen* graph becomes non-differentiable. To overcome this issue the method uses reinforcement learning method called policy gradients (Williams, 1992) to train the generator.

Classifier *clas*. *clas* is a softmax classifier that tries to determine whether a comment is inappropriate or not by processing only words from rationale (z, x) .

Joint learning. In order to learn to highlight inappropriate words from inappropriate comments we need *gen* and *clas* to cooperate. *gen* selects words and *clas* provides feedback on the quality of selected words. The feedback is based on the assumption that the words are selected correctly if *clas* is able to classify comment correctly based on the rationale (z, x) and vice versa.

Furthermore, there are some conditions on the rationale: it must to be short and meaningful (the selected words must be near each other) what is achieved by adding regularization controlled by hyperparameters λ_1 (which forces the rationales to have fewer words) and λ_2 (which forces the selected words to be in a row). The following loss function expresses these conditions:

$$\begin{aligned} \text{loss}(x, z, y') &= \|\text{clas}(z, x) - y'\|_2^2 \\ &+ \lambda_1 \|z\| + \lambda_2 \sum_{t=1}^{K-1} |z_t - z_{t+1}| \quad (1) \end{aligned}$$

where x is original comment text, z contains binary flags representing non/selection of each word

in x , (z, x) contains actual words selected to rationale, y' is correct output and K is the length of x and also z respectively.

From the loss function we can see that the training is based on a simple assumption that rationale is a subset of words that *clas* classifies correctly. If it is not classified correctly then the rationale is probably incorrect. This way we can learn to generate rationales without need to have word level annotations. We would like to make the point that training to generate these rationales is not done to improve the classification performance. It uses only the exact same data the classifier from Step 1 uses. Its only effect is to generate interpretable rationales behind the decisions classifier takes.

4 Experiments and Results

4.1 Dataset

We used a proprietary dataset of more than 20 million comments from a major Slovak news discussion platform. Over the years a team of moderators was considering reported comments and removing the inappropriate ones while also selecting a reason(s) from prepared list of possible discussion code violations. In this work we consider only those that were flagged because of following reasons: insults, racism, profanity or spam. The rest of the comments are considered appropriate. We split the dataset in train, validation and test set where validation and test set both were balanced to contain 10,000 appropriate and 10,000 inappropriate comments. Rest of the dataset forms the training set. Test and validation sets were sampled from the most recent months. During the training we balance it on batch level by supersampling inappropriate comments.

Highlights test set. We did not have any annotations on rationales in the dataset. We created a test set by manually selecting words that should form the rationales in randomly picked 100 comments. This way we created a test set containing 3,600 annotated words.

Word embeddings. We trained our own fast-Text embeddings (Bojanowski et al., 2017) on our dataset. These take into account character level information and are therefore suitable for inflected languages (such as Slovak) and online discussions where lots of grammatical and typing errors occur.

4.2 Inappropriate comments detection

We performed a hyperparameter grid search with our method. We experimented with different recurrent cells (RCNN and LSTM), depth (2, 3), hidden size (200, 300, 500), bi-directional RNN and in the case of RCNN also with cell order (2, 4). We also trained several non-neural methods for comparison. Results from this experiment are marked in Table 1. We measure accuracy as well as average precision (AP). The best results were achieved by bi-directional 2-layer RCNN with hidden size 300 and order 2. Deep neural network models outperform feature based models by almost 10% of accuracy. RCNN achieves results similar to LSTM but with approximately 8.5 times less parameters.

The results here might be significantly affected by noisy data. During the years many inappropriate comments went unnoticed and many appropriate comments were blocked if they were in inappropriate threads. Qualitative interviews we carried out with moderators indicate that our accuracy might be a bit higher.

We observed that model was the most confident about insulting and offensive comments. Thanks to sub-word based word embeddings the model can find profanities even when some characters within are replaced with numbers (e.g. *Insult* instead of *insult*) or there are arbitrary characters inserted into the word (e.g. *i..n..s..u..l..t* instead of *insult*).

To better understand the impact of this classifier we plotted its results using ROC curve in Figure 1. Here we can see how many comments a moderator needs to read to find certain percentile of inappropriate ones. E.g. when looking for 80% of inappropriate comments, only 20% of reviewed comments will be falsely flagged by the model.

4.3 Highlighting inappropriate parts

gen and *clas* are implemented as recurrent neural networks with RCNN cells. *gen* is a bi-directional 2-layer RCNN with hidden size equal 200 and order equal 2. Z-layer is realized as unidirectional RNN with hidden size equal 30. *clas* is an unidirectional 2-layer RCNN with hidden size equal 200 and order equal 2. For regularization hyperparameters we found values $\lambda_1 \in [5 \times 10^{-4}, 3 \times 10^{-3}]$ and $\lambda_2 \in [2\lambda_1, 4\lambda_1]$ to perform well.

We observed a significant instability during the training caused by formulation of our loss func-

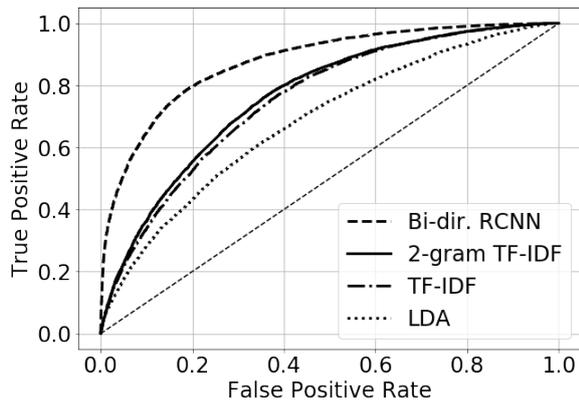


Figure 1: Receiver operating characteristic (ROC) of different models. We plot only one neural based solution as they almost completely overlap.

Model	# params	% acc	AUC
LDA (75 topics)	-	63.2	0.684
LSA (300 topics)	-	66.2	-
TF-IDF	-	68.8	0.754
2-gram TF-IDF	-	70.1	0.766
Unidir. RCNN	0.6M	79.0	0.870
Unidir. LSTM	3.3M	78.9	0.872
Bi-dir. RCNN	1.7M	79.4	0.872
Bi-dir. LSTM	14.6M	79.4	0.875

Table 1: Comparison of test set performance of multiple classification models. Baseline models (first four) use various text representations that are classified by boosted decision trees (1,000 trees) (Freund and Schapire, 1995).

tion. The model would often converge to a state where it would pick all the words or no words at all. Especially the cases when the model started to pick all the words proved to be impossible to overcome. In such cases we restarted the training from a different seed what increased the probability of a model converging successfully by a factor of five.

We evaluated following metrics of our models:

- *Precision* – how many of selected words were actually part of golden inappropriate data. Correct selection of words is a prerequisite for saving moderators’ time. Recall is not very important as we do not need to select all the inappropriate parts. One part is usually enough for the moderators to block a comment.
- *Rationale length* – the proportion of words selected into rationale. It is important to mea-

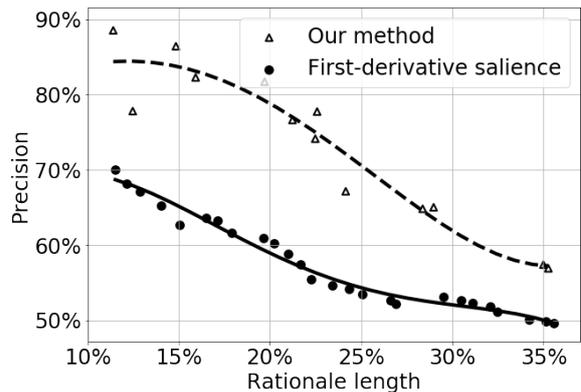


Figure 2: Test set performance of models highlighting inappropriate parts in comments.

sure this metric as we want our model to only pick a handful of strongly predictive words.

We compare our proposed model with the model based on first-derivative saliency (Li et al., 2016a). The comparison of models is shown in Figure 2. We can see that with length reduction the precision grows as expected. Our best models achieve a precision of nearly 90% while selecting 10–15% of words. We consider this to be very good result. Our method outperforms the saliency based one and also produces less scattered rationales. By this we mean that the average length of a segment of subsequently selected words is 2.5 for our method, but only 1.5 for saliency-based method. Instead of picking individual words our model tries to pick longer segments.

5 Conclusion

Moderating online discussions is time consuming error-prone activity. Major discussion platforms have millions of users so they need huge teams of moderators. We propose a method to speed up this process and make it more reliable. The novelty of our approach is in the application of a model interpretation method in this domain.

Instead of simply marking the comment as inappropriate, our method highlights the words that made the model think so. This is significant help for moderators as they can now read only small part of comment instead of its whole text. We believe that our method can significantly speed up the moderation process and user study is underway to confirm this hypothesis.

We evaluated our model on dataset from a major Slovak news discussion platform with more than

20 million comments. Results are encouraging as we were able to obtain good results on inappropriate comments detection task. We also obtained good results (nearly 90% precision) when highlighting inappropriate parts of these comments.

In the future we plan to improve the evaluation of highlighting. Instead of measuring the global precision, we plan to analyze how well it performs with various types of inappropriateness, such as racism, insults or spam. We are also looking into the possibility of incorporating additional data into our algorithm – other comments from the same thread, article for which the comments are created or even user profiles. These could help us improve our results or even detect the possibility of antisocial behavior before it even happens.

Acknowledgments

This work was partially supported by the Slovak Research and Development Agency under the contracts No. APVV-17-0267 - Automated Recognition of Antisocial Behaviour in Online Communities and No. APVV-15-0508 - Human Information Behavior in the Digital Space. The authors would like to thank for financial contribution from the Scientific Grant Agency of the Slovak Republic, grant No. VG 1/0646/15.

References

- B Thomas Adler, Luca De Alfaro, Santiago M Mola-Velasco, Paolo Rosso, and Andrew G West. 2011. Wikipedia vandalism detection: Combining natural language, metadata, and reputation features. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 277–288. Springer.
- Leila Arras, Franziska Horn, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. 2017. "what is relevant in a text document?": An interpretable machine learning approach. *PLOS ONE*, 12(8):1–23.
- Malika Aubakirova and Mohit Bansal. 2016. Interpreting neural networks to improve politeness comprehension. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2035–2041. Association for Computational Linguistics.
- Pinkesh Badjatiya, Shashank Gupta, Manish Gupta, and Vasudeva Varma. 2017. Deep learning for hate speech detection in tweets. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 759–760. International World Wide Web Conferences Steering Committee.
- Tao Lei Hrishikesh Joshi Regina Barzilay, Tommi Jaakkola, Katerina Tymoshenko, and Alessandro Moschitti Lluís Marquez. 2016. Semi-supervised question retrieval with gated convolutions. In *Proceedings of NAACL-HLT*, pages 1279–1289.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association of Computational Linguistics*, 5(1):135–146.
- Pete Burnap and Matthew L Williams. 2016. Us and them: identifying cyber hate on twitter across multiple protected characteristics. *EPJ Data Science*, 5(1):11.
- Justin Cheng, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2014. How community feedback shapes user behavior. In *Eighth International AAAI Conference on Weblogs and Social Media*.
- Justin Cheng, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2015. Antisocial behavior in online discussion communities. In *Ninth International AAAI Conference on Web and Social Media*.
- Nemanja Djuric, Jing Zhou, Robin Morris, Mihajlo Grbovic, Vladan Radosavljevic, and Narayan Bhamidipati. 2015. Hate speech detection with comment embeddings. In *Proceedings of the 24th international conference on world wide web*, pages 29–30. ACM.
- European court of human rights. 2015. Case of Delfi as v. Estonia. [https://hudoc.echr.coe.int/eng#{\"itemid\":\[\"001-155105\"\]}](https://hudoc.echr.coe.int/eng#{\). [Online; accessed February 23th, 2018].
- Yoav Freund and Robert E Schapire. 1995. A decision-theoretic generalization of on-line learning and an application to boosting. In *European conference on computational learning theory*, pages 23–37. Springer.
- Njagi Dennis Gitari, Zhang Zuping, Hanyurwimfura Damien, and Jun Long. 2015. A lexicon-based approach for hate speech detection. *International Journal of Multimedia and Ubiquitous Engineering*, 10(4):215–230.
- Andrej Karpathy, Justin Johnson, and Li Fei-Fei. 2015. Visualizing and understanding recurrent networks. *arXiv preprint arXiv:1506.02078*.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 107–117.
- Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. 2016a. Visualizing and understanding neural models in nlp. In *Proceedings of NAACL-HLT*, pages 681–691.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016b. Understanding neural networks through representation erasure. *CoRR*, abs/1612.08220.

- Yashar Mehdad and Joel Tetreault. 2016. Do characters abuse more than words? In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 299–303.
- W. James Murdoch, Peter J. Liu, and Bin Yu. 2018. Beyond word importance: Contextual decomposition to extract interactions from lstms. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive language detection in online user content. In *Proceedings of the 25th international conference on world wide web*, pages 145–153. International World Wide Web Conferences Steering Committee.
- John Pavlopoulos, Prodromos Malakasiotis, and Ion Androutsopoulos. 2017. Deeper attention to abusive user content moderation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1125–1135. Association for Computational Linguistics.
- Manoel Horta Ribeiro, Pedro H Calais, Yuri A Santos, Virgílio AF Almeida, and Wagner Meira Jr. 2018a. Characterizing and detecting hateful users on twitter. *arXiv preprint arXiv:1803.08977*.
- Manoel Horta Ribeiro, Pedro H Calais, Yuri A Santos, Virgílio AF Almeida, and Wagner Meira Jr. 2018b. “like sheep among wolves”: Characterizing hateful users on twitter.
- Anna Schmidt and Michael Wiegand. 2017. 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.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pages 5–32. Springer.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J. Smola, and Eduard H. Hovy. 2016. Hierarchical attention networks for document classification. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 1480–1489.
- Matthew D. Zeiler and Rob Fergus. 2014. Visualizing and understanding convolutional networks. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I*, pages 818–833.
- Haoti Zhong, Hao Li, Anna Squicciarini, Sarah Rajtmajer, Christopher Griffin, David Miller, and Cornelia Caragea. 2016. Content-driven detection of cyberbullying on the instagram social network. In *Proceedings of the Twenty-Fifth International Joint*

Conference on Artificial Intelligence, pages 3952–3958. AAAI Press.