

WLV-RIT at SemEval-2021 Task 5: A Neural Transformer Framework for Detecting Toxic Spans

Tharindu Ranasinghe¹, Diptanu Sarkar², Marcos Zampieri², Alexander Ororbia²

¹University of Wolverhampton, UK

²Rochester Institute of Technology, USA

T.D.RanasingheHettiarachchige@wlv.ac.uk

Abstract

In recent years, the widespread use of social media has led to an increase in the generation of toxic and offensive content on online platforms. In response, social media platforms have worked on developing automatic detection methods and employing human moderators to cope with this deluge of offensive content. While various state-of-the-art statistical models have been applied to detect toxic posts, there are only a few studies that focus on detecting the words or expressions that make a post offensive. This motivates the organization of the SemEval-2021 Task 5: Toxic Spans Detection competition, which has provided participants with a dataset containing toxic spans annotation in English posts. In this paper, we present the WLV-RIT entry for the SemEval-2021 Task 5. Our best performing neural transformer model achieves an 0.68 F1-Score. Furthermore, we develop an open-source framework for multilingual detection of offensive spans, i.e., MUDES, based on neural transformers that detect toxic spans in texts.

1 Introduction

The widespread adoption and use of social media has led to a drastic increase in the generation of abusive and profane content on the web. To counter this deluge of negative content, social media companies and government institutions have turned to developing and applying computational models that can identify the various forms of offensive content online such as aggression (Kumar et al., 2018, 2020), cyber-bullying (Rosa et al., 2019), and hate speech (Ridenhour et al., 2020). Prior work has either designed methods for identifying conversations that are likely to go awry (Zhang

et al., 2018; Chang et al., 2020) or detecting offensive content and labelling posts at the instances level – this has been the focus in the recent shared tasks like HASOC at FIRE 2019 (Mandl et al., 2019a) and FIRE 2020 (Mandl et al., 2020), GermEval 2019 Task 2 (Struß et al., 2019), TRAC (Kumar et al., 2018, 2020), HatEval (Basile et al., 2019a), OffensEval at SemEval-2019 (Zampieri et al., 2019b) and SemEval-2020 (Zampieri et al., 2020).

With respect to identifying offensive language in conversations, comments, and posts, noticeable progress has been made with a variety of large, annotated datasets made available in recent years (Pitenis et al., 2020; Rosenthal et al., 2020). The identification of the particular text spans that make a post offensive, however, has been mostly neglected (Mathew et al., 2021) as current state-of-the-art offensive language identification models flag the entire post or comment but do not actually highlight the offensive parts. The pressing need for toxic span detection models to assist human content moderation, processing and flagging content in a more interpretable fashion, has motivated the organization of the SemEval-2021 Task 5: Toxic Spans Detection (Pavlopoulos et al., 2021).

In this paper, we present the WLV-RIT submission to the SemEval-2021 Task 5. We explore several statistical learning models and report the performance of the best model, which is based on a neural transformer. Next, we generalise our approach to an open-source framework called MUDES: Multilingual Detection of Offensive Spans (Ranasinghe and Zampieri, 2021a). Alongside the framework, we also release the pre-trained models as well as a user-friendly web-based User Interface (UI) based on Docker, which provides the functionality of automatically identifying the offensive spans in a given input text.

WARNING: This paper contains text excerpts and words that are offensive in nature.

2 Related Work

Datasets Over the past several years, multiple post-level, offensive language benchmark datasets have been released. In [Zampieri et al. \(2019a\)](#), the authors compiled an offensive language identification dataset with a three-layer hierarchical annotation scheme – profanity, category, and target identification. [Rosenthal et al. \(2020\)](#) further extended the dataset using a semi-supervised model that was trained with over nine million annotated English tweets. Recently, [Mathew et al. \(2021\)](#) released the first benchmark dataset which covered the three primary areas of online hate-speech detection. The dataset contained a 3-class classification problem (hate-speech, offensive, or neither), a targeted community, as well as the spans that make the text hateful or offensive. Furthermore, offensive language datasets have been annotated in other languages such as Arabic ([Mubarak et al., 2017](#)), Danish ([Sigurbergsson and Derczynski, 2020](#)), Dutch ([Tulkens et al., 2016](#)), French ([Chiril et al., 2019](#)), Greek ([Pitenis et al., 2020](#)), Portuguese ([Fortuna et al., 2019](#)), Spanish ([Basile et al., 2019b](#)), and Turkish ([Çöltekin, 2020](#)).

Apart from the dataset released for SemEval-2021 Task 5, HateXplain ([Mathew et al., 2021](#)) is, to the best of our knowledge, the only dataset that we could find that has been annotated at the word level. The dataset consists of 20,000 posts from Gab and Twitter. Each data sample is annotated with one of the hate/offensive/normal labels, communities being targeted, and words of the text are marked by the annotators who support the label.

Models In the past, trolling, aggression, and cyberbullying identification tasks on social media data have been approached using machine and deep learning-focused models ([Kumar et al., 2018](#)). Across several studies ([Malmasi and Zampieri, 2017, 2018](#); [Waseem and Hovy, 2016](#)) researchers have noted that n -gram based features are very useful when building reliable, automated hate-speech detection models. Statistical learning models aided with natural language processing (NLP) techniques are frequently used for post-level offensive and hateful language detection ([Davidson et al., 2017](#); [Indurthi et al., 2019](#)). Given the increased use of deep learning in NLP tasks, offensive language identification has seen the introduction of methods based on convolutional neural networks (CNNs) and Long Short-term Memory

(LSTM) networks ([Badjatiya et al., 2017](#); [Gambäck and Sikdar, 2017](#); [Hettiarachchi and Ranasinghe, 2019](#)). The most common approach has been to use a word/character embedding model such as Word2vec ([Mikolov et al., 2013](#)), GloVe ([Pennington et al., 2014](#)), or fastText ([Mikolov et al., 2018](#)) to embed words/tokens and then feed them to an artificial neural network (ANN) ([Zampieri et al., 2019b](#)).

With the introduction of BERT ([Devlin et al., 2019](#)), neural transformer models have become popular in offensive language identification. In hate speech and offensive content identification in Indo-European languages, the BERT model has been shown to outperform GRU (Gated Recurrent Unit) and LSTM-based models ([Ranasinghe et al., 2019](#)). In [Mandl et al. \(2019b\)](#), the best performing teams on the task employed BERT-based pre-trained models that identified the type of hate and target of a (text) post.

The SemEval-2019 Task 6 ([Zampieri et al., 2019b](#)) presented the challenge of identifying and categorizing offensive posts on social media, which included three sub-tasks. In sub-task A: offensive language identification, [Liu et al. \(2019a\)](#) applied a pre-trained BERT model to achieve the highest F1 score. In Sub-task B: automatic categorization of offense types, BERT-based models also achieved competitive rankings. We noticed similar trends in SemEval-2020 Task 12 ([Zampieri et al., 2020](#)) as well. Not limited to English, transformer models have yielded strong results in resource-scarce languages like Bengali ([Ranasinghe and Zampieri, 2020](#)) and Malayalam ([Ranasinghe et al., 2020](#)) along with cross-lingual transfer learning from resource-rich languages ([Ranasinghe and Zampieri, 2020, 2021b](#)). Nonetheless, despite the recent success of statistical learning in offensive language detection problems, due to the lack of finer-grained, detailed datasets, models are limited in their ability to predict word-level labels.

3 Task and Dataset

In the SemEval-2021 Task 5 dataset, the sequence of words that makes a particular post or comment toxic is defined as a *toxic span*. The dataset for this task is extracted from posts in the Civil Comments Dataset that have been found to be toxic. The practice dataset has 690 instances out of which 43 instances do not contain any toxic spans. The training dataset has a total of 7,939 instances and

Post	Offensive Spans
Stupid hatcheries have completely fucked everything	[0, 1, 2, 3, 4, 5, 34, 35, 36, 37, 38, 39]
Victimitis: You are such an asshole .	[28, 29, 30, 31, 32, 33, 34]
So is his mother. They are silver spoon parasites.	[]
You're just silly .	[12, 13, 14, 15, 16]

Table 1: Four comments from the dataset along with their annotations. The offensive words are displayed in red and the spans are indicated by the character position in the instance.

comprises 485 instances without any toxic spans. Each instance is composed of a list of toxic spans and the post (in English). In Table 1, we present four randomly selected examples from the training dataset along with their annotations.




Figure 1: The Bi-LSTM-CRF model. Green squares represent the top CRF layer. Non-offensive and offensive tokens are shown as 0 and 1, respectively.

4 Methodology

4.1 Lexicon-based Word Match

Lexicon-based word-matching algorithms often achieve balanced results. For the lexicon, we collected profanity words from online resources^{1,2}. Then, we added the toxic words present in the training dataset and we run a simple word matching algorithm the trie data structure. As anticipated, the algorithm does not evaluate the toxic spans contextually and misses censored swear words. For instance, the word $f^{**}k$ is missed, which is not present in the lexicon. Nonetheless, this result provides as a useful baseline performance measurement for the task.

4.2 Recurrent Networks: Long Short-Term Memory

Long Short-term Memory (LSTM) is a recurrent neural network model that uses feedback connec-

¹<https://www.cs.cmu.edu/~biglou/resources/bad-words.txt>

²<https://github.com/RobertJGabriel/Google-profanity-words>

tions to model temporal dependencies (past-to-present) in sequential data. Bidirectional LSTM (Bi-LSTM) is capable of learning contextual information both forwards and backwards in time compared to conventional LSTMs. In this study, we used the Bi-LSTM architecture given this bidirectional ability to model temporal dependencies. Conditional random fields (CRF) (Lafferty et al., 2001) are a statistical model that are capable of incorporating context information and are highly used for sequence labeling tasks. A CRF connected to the top of the Bi-LSTM model provides a powerful way to model relationships between consecutive outputs (across time) and provides a means to efficiently utilize past and future tag information to predict the current tag.

The final hybrid model is comparable to the previous state-of-the-art sequence tagging Bi-LSTM-CRF model (Huang et al., 2015). Figure 1 presents the Bi-LSTM-CRF architecture we designed for this study, which has 4.2 million trainable parameters. We trained the model on mini-batches of 16 samples with a 0.005 learning rate for 5 epochs with a maximum sequence length of 200.

4.3 Neural Transformers

Recently, pre-trained language models have been shown to be quite useful across a variety of NLP tasks, particularly those based on bidirectional neural transformers such as BERT (Devlin et al., 2019; Li et al., 2019). Transformer-based models have also been shown to be highly effective in sequence classification tasks such as named entity recognition (NER) (Luoma and Pyysalo, 2020). In our work, we extend the BERT model by integrating a token level classifier. The token-level classifier is a linear transformation that takes the last hidden state of the sequence as the input and produces a label for each token as its output. In this case, each token will be predicted to have one of two possible labels – toxic or not toxic. We fine-tuned the uncased BERT transformer model with a maximum

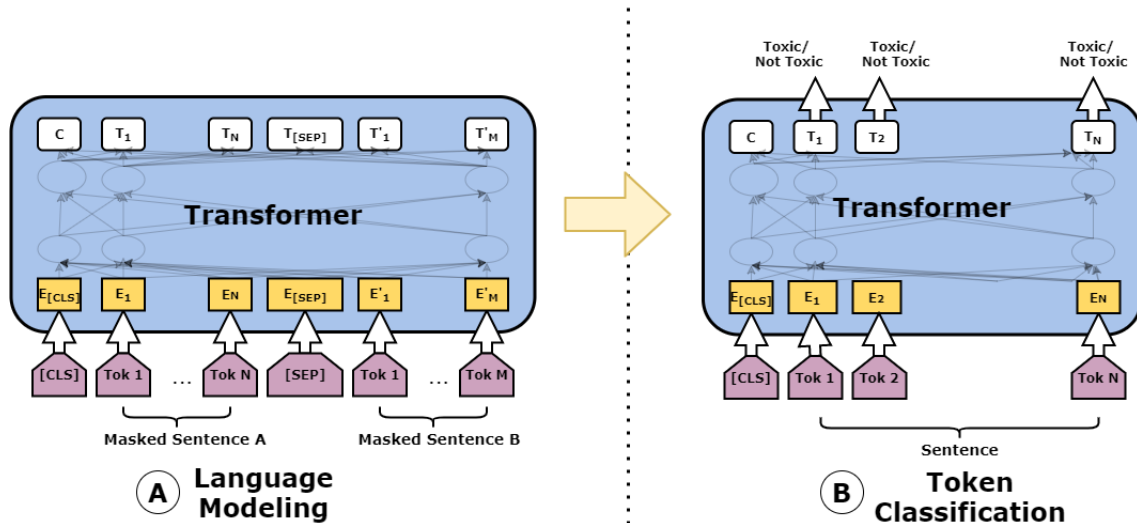


Figure 2: The two-part model architecture. Part A depicts the language model and Part B is the token classifier. (Ranasinghe and Zampieri, 2021a)

sequence length of 400 with batches of size of 16.

We also experimented with customising the layers in between the BERT transformer and token-classification layer by adding a CRF layer between them given that it has been shown that BERT-CRF architectures often outperform BERT baselines in similar sequence labeling tasks (Huang et al., 2019; Souza et al., 2020). Therefore, we added a sequential CRF layer on top of the BERT transformer and further incorporated dropout (probability of dropping a neuron was 0.2) to introduce some regularization. Unfortunately, in our experiments, we found that adding a CRF layer does not significantly improve the final generalization results. Additionally, we experimented with transfer learning to identify if a further boost in model generalization was possible if we first trained a basic BERT transformer on HateXplain (Mathew et al., 2021) and then fine-tuned it using our extended architecture as described above. However, the transfer learning process did not improve results any further.

Development of MUDES Given the success we observed using neural transformers such as BERT, we developed a (software) framework we call MUDES (Ranasinghe and Zampieri, 2021a): Multilingual Detection of Offensive Spans, an open-source framework based on transformers to detect toxic spans in texts. MUDES offers several capabilities in addition to the (automatic) token classification we described earlier. MUDES has the following components: *a) Language Modeler*: Fine-tuning transformer models using masked language

modeling before performing the downstream task often leads to better results (Ranasinghe and Hettiarachchi, 2020) and MUDES incorporates this, *b) Transformer Type Variety*: since there are many varieties of neural transformers, e.g., XLNet (Yang et al., 2019), RoBERTa (Liu et al., 2019b) that have been shown to outperform BERT-based architectures (Ranasinghe and Hettiarachchi, 2020; Hettiarachchi and Ranasinghe, 2020a), our software framework provides support for these architectures, and, finally, *c) Model Ensembling*: multiple MUDES models with different random seeds can be trained and the final model prediction is the majority vote from all the models, aligning with the approach taken in Hettiarachchi and Ranasinghe (2020b, 2021); Jauhainen et al. (2021).

The complete architecture of MUDES is depicted in Figure 2. We used several popular transformer models including BERT (Devlin et al., 2019), XLNET (Yang et al., 2019), RoBERTa (Liu et al., 2019b), SpanBERT (Joshi et al., 2020), and ALBERT (Lan et al., 2020). We compared these transformer architectures against the spaCy token classifier baseline (reported by the competition organisers) and report these results in Section 5. Since adding a CRF layer did not improve the results in our models, we do not add this to MUDES.

Parameter optimization involved mini-batches of 8 samples using the Adam update rule (global learning rate was $2e-5$ and a linear warm-up schedule over 10% of the training data was used). Models were evaluated using a validation subset that contained 20% of the training data. Early stopping

was executed if the validation loss did not improve over 10 evaluation steps. Models were trained for 3 epochs on an Nvidia Tesla K80 GPU using only the training set provided.

5 Evaluation and Results

For evaluation, we followed the same procedure that the task organisers have used to evaluate the systems.

Let system A_i return a set $S_{A_i}^t$ of character offsets for parts of a text post that have been found to be toxic. Let G_t be the character offsets of the ground truth annotations of t . We compute the F1 score of system A_i with respect to the ground truth G for post t as mentioned in Equation 1 where $|\cdot|$ denotes set cardinality. P^t and R^t measure the precision and recall, respectively.

$$F_1^t(A_i, G) = \frac{2 \cdot P^t(A_i, G) \cdot R^t(A_i, G)}{P^t(A_i, G) + R^t(A_i, G)} \quad (1)$$

Model	Trial F1	Test F1
MUDES RoBERTa	0.6886	0.6801
MUDES BERT	0.6771	0.6698
MUDES SPANBert	0.6751	0.6675
MUDES XLNet	0.6722	0.6653
BERT	0.6738	0.6538
BERT-CRF	0.6643	0.6517
BERT HateXplain	0.6387	0.6326
spaCy baseline	0.5976	0.5976
Bi-LSTM-CRF	0.5631	0.5398
Lexicon word match	0.3378	0.4086

Table 2: Results ordered by test F1 score. The Trial F1 column shows the F1 scores on the trial set and the Test F1 column shows the F1 scores for test set.

Observe in Table 2 that all of our deep neural-based models outperformed the spaCy baseline while the lexicon-based word match algorithm provided fairly good results despite it being an unsupervised method. Our best model is the MUDES RoBERTa model which scored 0.68 F1 score in the test set and is very compatible with the 0.70 F1 score that the best model scored in the competition. Furthermore, it is clear that the additional features supported by our MUDES framework, e.g., language modeling and ensembling, improves the results over a vanilla BERT transformer.

6 Conclusion and Future Work

In this paper, we presented the WLV-RIT approach for tackling the SemEval-2021 Task 5: Toxic Spans Detection. SemEval-2021 Task 5 provided participants with the opportunity of testing computational models to identify token spans in toxic posts as opposed to previous related SemEval tasks such as HatEval and OffensEval that provided participants with datasets annotated at the instance level. We believe that word-level predictions are an important step towards explainable offensive language identification.

We experimented with several methods including a lexicon-based word match, LSTMs, and neural transformers. Our results demonstrated that transformer models offered the best generalization results and, given the success observed, we developed MUDES, an open-source software framework based on neural transformers focused on detecting toxic spans in texts. With MUDES, we release two English models that performed best for this task (Ranasinghe and Zampieri, 2021a). A large model; en-large based on roberta-large which is more accurate, but has a low efficiency regarding space and time. The base model based on xlnet-base-cased; en-base is efficient, but has a comparatively low accuracy than the en-large model. All pre-trained models are available on Hugging Face Model Hub (Wolf et al., 2020)³. We also make MUDES available as a Python package⁴ and set up as an open-source project⁵. In addition, a prototype User Interface (UI) of MUDES has been made accessible to the general public⁶ based on Docker⁷.

In terms of future work, we would like to experiment with multi-task (neural) architectures that can be used for offensive language identification capable of carrying out predictions at both the word-level and post-level jointly. Furthermore, we would like to evaluate multi-task architectures on multi-domain and multilingual settings as well as broaden our experimental comparison to other types of recurrent network models, such as the Delta-RNN (Ororbia II et al., 2017).

³Available on <https://huggingface.co/mudes>

⁴Available at <https://pypi.org/project/mudes/>

⁵The MUDES GitHub repository is available at <https://github.com/tharindudr/MUDES>

⁶The UI can be accessed from <http://rgcl.wlv.ac.uk/mudes/>

⁷Available at <https://hub.docker.com/r/tharindudr/mudes>

Acknowledgments

We would like to thank the shared task organizers for making this interesting dataset available. We further thank the anonymous SemEval reviewers for their insightful feedback.

References

- Pinkesh Badjatiya, Shashank Gupta, Manish Gupta, and Vasudeva Varma. 2017. Deep learning for hate speech detection in tweets. In *Proceedings of WWW*.
- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019a. Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter. In *Proceedings of SemEval*.
- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019b. SemEval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter. In *Proceedings of SemEval*.
- Çağrı Çöltekin. 2020. A Corpus of Turkish Offensive Language on Social Media. In *Proceedings of LREC*.
- Jonathan P Chang, Justin Cheng, and Cristian Danescu-Niculescu-Mizil. 2020. Don't let me be misunderstood: Comparing intentions and perceptions in online discussions. In *Proceedings of WWW*.
- Patricia Chiril, Farah Benamara Zitoune, Véronique Moriceau, Marlène Coulomb-Gully, and Abhishek Kumar. 2019. Multilingual and multitarget hate speech detection in tweets. In *Proceedings of TALN*.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated Hate Speech Detection and the Problem of Offensive Language. In *Proceedings of ICWSM*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL*.
- Paula Fortuna, Joao Rocha da Silva, Leo Wanner, Sérgio Nunes, et al. 2019. A Hierarchically-labeled Portuguese Hate Speech Dataset. In *Proceedings of ALW*.
- Björn Gambäck and Utpal Kumar Sikdar. 2017. Using convolutional neural networks to classify hate-speech. In *Proceedings of ALW*.
- Hansi Hettiarachchi and Tharindu Ranasinghe. 2019. Emoji powered capsule network to detect type and target of offensive posts in social media. In *Proceedings of RANLP*.
- Hansi Hettiarachchi and Tharindu Ranasinghe. 2020a. BRUMS at SemEval-2020 task 3: Contextualised embeddings for predicting the (graded) effect of context in word similarity. In *Proceedings of SemEval*.
- Hansi Hettiarachchi and Tharindu Ranasinghe. 2020b. InfoMiner at WNUT-2020 task 2: Transformer-based covid-19 informative tweet extraction. In *Proceedings of W-NUT*.
- Hansi Hettiarachchi and Tharindu Ranasinghe. 2021. TransWiC at SemEval-2021 Task 2: Transformer-based Multilingual and Cross-lingual Word-in-Context Disambiguation. In *Proceedings of SemEval*.
- Weipeng Huang, Xingyi Cheng, Taifeng Wang, and Wei Chu. 2019. BERT-Based Multi-Head Selection for Joint Entity-Relation Extraction. In *Proceedings of NLPCC*.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Vijayasaradhi Indurthi, Bakhtiyar Syed, Manish Shrivastava, Nikhil Chakravartula, Manish Gupta, and Vasudeva Varma. 2019. FERMI at SemEval-2019 task 5: Using sentence embeddings to identify hate speech against immigrants and women in Twitter. In *Proceedings of SemEval*.
- Tommi Jauhiainen, Tharindu Ranasinghe, and Marcos Zampieri. 2021. Comparing approaches to dravidian language identification. In *Proceedings of VarDial*.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving pre-training by representing and predicting spans. *Proceedings of TACL*.
- Ritesh Kumar, Atul Kr Ojha, Shervin Malmasi, and Marcos Zampieri. 2018. Benchmarking aggression identification in social media. In *Proceedings of TRAC*.
- Ritesh Kumar, Atul Kr. Ojha, Shervin Malmasi, and Marcos Zampieri. 2020. Evaluating Aggression Identification in Social Media. In *Proceedings of TRAC*.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of ICML*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. In *Proceedings of ICLR*.
- W. Li, S. Gao, H. Zhou, Z. Huang, K. Zhang, and W. Li. 2019. The Automatic Text Classification Method Based on BERT and Feature Union. In *Proceedings of ICPADS*.

- Ping Liu, Wen Li, and Liang Zou. 2019a. NULI at SemEval-2019 task 6: Transfer learning for offensive language detection using bidirectional transformers. In *Proceedings of SemEval*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.
- Jouni Luoma and Sampo Pyysalo. 2020. Exploring Cross-sentence Contexts for Named Entity Recognition with BERT. In *Proceedings of COLING*.
- Shervin Malmasi and Marcos Zampieri. 2017. Detecting Hate Speech in Social Media. In *Proceedings of RANLP*.
- Shervin Malmasi and Marcos Zampieri. 2018. Challenges in Discriminating Profanity from Hate Speech. *Journal of Experimental & Theoretical Artificial Intelligence*, 30:1–16.
- Thomas Mandl, Sandip Modha, Anand Kumar M, and Bharathi Raja Chakravarthi. 2020. Overview of the hasoc track at fire 2020: Hate speech and offensive language identification in tamil, malayalam, hindi, english and german. In *Proceedings of FIRE*.
- Thomas Mandl, Sandip Modha, Prasenjit Majumder, Daksh Patel, Mohana Dave, Chintak Mandlia, and Aditya Patel. 2019a. Overview of the hasoc track at fire 2019: Hate speech and offensive content identification in indo-european languages. In *Proceedings of FIRE*.
- Thomas Mandl, Sandip Modha, Prasenjit Majumder, Daksh Patel, Mohana Dave, Chintak Mandlia, and Aditya Patel. 2019b. Overview of the hasoc track at fire 2019: Hate speech and offensive content identification in indo-european languages. In *Proceedings of FIRE*.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2021. Hatexplain: A benchmark dataset for explainable hate speech detection. In *Proceedings of AACL*.
- Tomas Mikolov, Edouard Grave, Piotr Bojanowski, Christian Puhersch, and Armand Joulin. 2018. Advances in pre-training distributed word representations. In *Proceedings of LREC*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of NeurIPS*.
- Hamdy Mubarak, Kareem Darwish, and Walid Magdy. 2017. Abusive language detection on Arabic social media. In *Proceedings of ALW*.
- Alexander G Ororbia II, Tomas Mikolov, and David Reitter. 2017. Learning simpler language models with the differential state framework. *Neural computation*, 29(12):3327–3352.
- John Pavlopoulos, Léo Laugier, Jeffrey Sorensen, and Ion Androutsopoulos. 2021. Semeval-2021 task 5: Toxic spans detection. In *Proceedings of SemEval*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of EMNLP*.
- Zeses Pitenis, Marcos Zampieri, and Tharindu Ranasinghe. 2020. Offensive Language Identification in Greek. In *Proceedings of LREC*.
- Tharindu Ranasinghe, Sarthak Gupte, Marcos Zampieri, and Ifeoma Nwogu. 2020. WLV-RIT at HASOC-Dravidian-CodeMix-FIRE2020: Offensive Language Identification in Code-switched YouTube Comments. In *Proceedings of FIRE*.
- Tharindu Ranasinghe and Hansi Hettiarachchi. 2020. BRUMS at SemEval-2020 task 12: Transformer based multilingual offensive language identification in social media. In *Proceedings of SemEval*.
- Tharindu Ranasinghe and Marcos Zampieri. 2020. Multilingual Offensive Language Identification with Cross-lingual Embeddings. In *Proceedings of EMNLP*.
- Tharindu Ranasinghe and Marcos Zampieri. 2021a. MUDES: Multilingual Detection of Offensive Spans. In *Proceedings of NAACL*.
- Tharindu Ranasinghe and Marcos Zampieri. 2021b. Multilingual Offensive Language Identification for Low-resource Languages. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)*.
- Tharindu Ranasinghe, Marcos Zampieri, and Hansi Hettiarachchi. 2019. BRUMS at HASOC 2019: Deep Learning Models for Multilingual Hate Speech and Offensive Language Identification. In *Proceedings of FIRE*.
- Michael Ridenhour, Arunkumar Bagavathi, Elaheh Raisi, and Siddharth Krishnan. 2020. Detecting Online Hate Speech: Approaches Using Weak Supervision and Network Embedding Models. In *Proceedings of SBP-BRIMS*.
- Hugo Rosa, N Pereira, Ricardo Ribeiro, Paula Costa Ferreira, Joao Paulo Carvalho, S Oliveira, Luísa Coheur, Paula Paulino, AM Veiga Simão, and Isabel Trancoso. 2019. Automatic cyberbullying detection: A systematic review. *Computers in Human Behavior*, 93:333–345.
- Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Marcos Zampieri, and Preslav Nakov. 2020. A large-scale semi-supervised dataset for offensive language identification. *arXiv preprint arXiv:2004.14454*.

- Gudbjartur Ingi Sigurbergsson and Leon Derczynski. 2020. Offensive Language and Hate Speech Detection for Danish. In *Proceedings of LREC*.
- Fábio Souza, Rodrigo Nogueira, and Roberto Lotufo. 2020. Portuguese Named Entity Recognition using BERT-CRF. *arXiv preprint arXiv:1909.10649*.
- Julia Maria Struß, Melanie Siegel, Josef Ruppenhofer, Michael Wiegand, Leibniz ScienceCampus, and Manfred Klenner. 2019. Overview of germeval task 2, 2019 shared task on the identification of offensive language. In *Proceedings of KONVENS*.
- Stéphan Tulkens, Lisa Hilte, Elise Lodewyckx, Ben Verhoeven, and Walter Daelemans. 2016. A Dictionary-based Approach to Racism Detection in Dutch Social Media. In *Proceedings of TA-COS*.
- Zeerak Waseem and Dirk Hovy. 2016. Hateful symbols or hateful people? predictive features for hate speech detection on Twitter. In *Proceedings of NAACL Student Research Workshop*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of EMNLP*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. XLNet: Generalized Autoregressive Pretraining for Language Understanding. In *Proceedings of NeurIPS*.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019a. Predicting the type and target of offensive posts in social media. In *Proceedings of NAACL*.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019b. SemEval-2019 Task 6: Identifying and Categorizing Offensive Language in Social Media (OffenseEval). In *Proceedings of SemEval*.
- Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. SemEval-2020 Task 12: Multilingual Offensive Language Identification in Social Media (OffenseEval 2020). *Proceedings of SemEval*.
- Justine Zhang, Jonathan Chang, Cristian Danescu-Niculescu-Mizil, Lucas Dixon, Yiqing Hua, Dario Taraborelli, and Nithum Thain. 2018. Conversations gone awry: Detecting early signs of conversational failure. In *Proceedings of ACL*.