

ur-iw-hnt at GermEval 2021: An Ensembling Strategy with Multiple BERT Models

Hoi Nam Tran

Information Science
University of Regensburg
Regensburg, Germany

Hoi-Nam.Tran@student.ur.de

Udo Kruschwitz

Information Science
University of Regensburg
Regensburg, Germany

Udo.Kruschwitz@ur.de

Abstract

This paper describes our approach (ur-iw-hnt) for the Shared Task of GermEval2021 to identify toxic, engaging, and fact-claiming comments. We submitted three runs using an ensembling strategy by majority (hard) voting with multiple different BERT models of three different types: German-based, Twitter-based, and multilingual models. All ensemble models outperform single models, while BERTtweet is the winner of all individual models in every subtask. Twitter-based models perform better than GermanBERT models, and multilingual models perform worse but by a small margin.

1 Introduction

Moderation of popular social media networks is a difficult task. Facebook alone has almost 2.8 billion active users on April 2021 (Kemp, 2021). Moderating discussions between users simultaneously all day is an impossible task, so moderators need help with this work. Also, fully automated solutions for content moderation are not possible, and human input is still required (Cambridge Consultants, 2019). An AI-based helper solution for harmful content detection is needed to make social networking less toxic and more pleasant instead.

The Shared Task of GermEval2021 focuses on highly relevant topics for moderators and community managers to moderate online discussion platforms (Risch et al., 2021). The challenge is not to specialize in one broad NLP task like harmful content detection but to detect other essential categories like which comments are engaging or fact-claiming.

We participated in all three subtasks (toxic, engaging and fact-claiming comment classification) to test our ensemble model to see whether multiple BERT-based models provide robust performance for different tasks without further customization. Moderators would benefit from a working system

without having to change models or settings all the time.

This report discusses in detail the three runs we submitted in the GermEval2021 Shared Task (Risch et al., 2021). We start with a brief reflection on related work, only focussing on aspects that are closely aligned with the subtasks. We then explain the dataset and the shared tasks in more detail. Next, we present our experiments, some discussions of the results, and we finally draw some conclusions.

To encourage reproducibility of experimental work, we make all code available via GitHub¹.

2 Related Work

Detecting harmful content in social media platforms is not only a monolingual but a multilingual issue. A multilingual toxic text detection classifier uses a fusion strategy employing mBERT and XLM-RoBERTa on imbalanced sample distributions (Song et al., 2021). Deep learning ensembles also show their effectiveness in hate speech detection (Zimmerman et al., 2019). A taxonomy of engaging comments contains different possible classifications (Risch and Krestel, 2020). With the increasing spread of misinformation, more collaborations with IT companies specialized in fact-checking and more intelligence and monitoring tools are available to help to identify harmful content (Arnold, 2020). An attempt to fully automate fact-checking is the tool called ClaimBuster (Hassan et al., 2015). Another tool named CrowdTangle monitors social media platforms and alerts the user if specific keywords are triggered so manual fact-claim checking can be done (Arnold, 2020). In addition, an annotation schema for claim detection is also available (Konstantinovskiy et al., 2021).

¹<https://github.com/HN-Tran/GermEval2021>

Dataset	Label	Subtask 1	Subtask 2	Subtask 3
Training	0	2122	2379	2141
	1	1122	865	1103
Test	0	594	691	630
	1	350	253	314

Table 1: Provided training and test dataset

3 Dataset & Shared Task

The dataset for the Shared Task of GermEval2021 consists of 3,244 annotated user discussion comments from a Facebook page of the German news broadcast in the timeframe of February to July 2019, labeled by four annotators in three different categories for binary classification: Toxic comments, engaging comments and fact-claiming comments (Risch et al., 2021). Since the labels are imbalanced, we first applied a stratified split onto the dataset so that 80% is for training. We then again apply a stratified split on what is left into two halves, the first part is the development set, and the second part is the holdout set for evaluation which we call the evaluation set here. After training, the ensemble strategy predicts the test dataset, consisting of 944 comments. Table 1 shows the imbalance in favor of the negative label. The organizers of GermEval2021 chose the metric Krippendorff’s alpha to check each task’s intercoder reliability (Risch et al., 2021).

3.1 Toxic Comment Classification

Toxic comments include many harmful and dangerous offenses like ”hate speech, insults, threats, vulgar advertisements and misconceptions about political and religious tendencies” (Song et al., 2021). Such behavior only leads to users leaving the discussion or manual bans by the moderator, which can be overwhelming depending on the number of active toxic users (Risch and Krestel, 2020). For this subtask, the annotator agreement in the usage of insults, vulgar and sarcastic language is $0.73 < \alpha < 0.89$, and in the discrimination, discreditation, accusations of lying or threats of violence, the agreement is at $0.83 < \alpha < 0.90$ (Risch et al., 2021).

3.2 Engaging Comment Classification

Engaging comments are, in general, attractive for users to participate in online discussions and get more interactions with other online users in the form of replies and upvotes. A taxonomy of engaging comments has been proposed to identify

these comments for detection and classification, so moderators and community managers can reward these comments or posts (Risch and Krestel, 2020). This task has three different categories (Risch et al., 2021):

- Juristification, solution proposals, sharing of personal experiences ($0.71 < \alpha < 0.89$)
- Empathy with regard to other users’ standpoints ($0.79 < \alpha < 0.91$)
- Polite interaction, mutual respect, mediation ($0.85 < \alpha < 1$)

3.3 Fact-Claiming Comment Classification

Detecting factual claims is part of the fact-checking process (Konstantinovskiy et al., 2021; Babakar and Moy, 2016; Nakov et al., 2021). The challenge here is to identify claims that have not been fact-checked before and go beyond one sentence that fits into this subtask (Babakar and Moy, 2016). Annotator’s agreement in fact assertion and evidence provision is at $0.73 < \alpha < 0.84$ (Risch et al., 2021).

4 Experiments

4.1 System Architecture

For our system architecture (see Figure 1), we use three Python libraries/tools. Deep-Translator² translates all the German comments into English by choosing an external service, in our case, the free public Google Translate service. We use two different libraries for classification: Ernie³ and Simple Transformers⁴. Both work on different versions of HuggingFace’s Transformers (Wolf et al., 2020) and thus differently: Ernie is a beginner-friendly library last updated in 2020, based on Keras / TensorFlow 2, and uses the optimizer Adam (Kingma and Ba, 2015). Simple Transformers is based on PyTorch and has more extensive options for hyperparameter tuning and training customizations with AdamW (Loshchilov and Hutter, 2019) as the default optimizer. The default hyperparameter values for our experiments, as recommended for BERT, are in Table 2. The only pre-processing step is the tokenization by each BERT model using these libraries. Because of time constraints, cross-validation has not been conducted. After training

²<https://github.com/nidhaloff/deep-translator>

³<https://github.com/labteral/ernie>

⁴<https://simpletransformers.ai/>

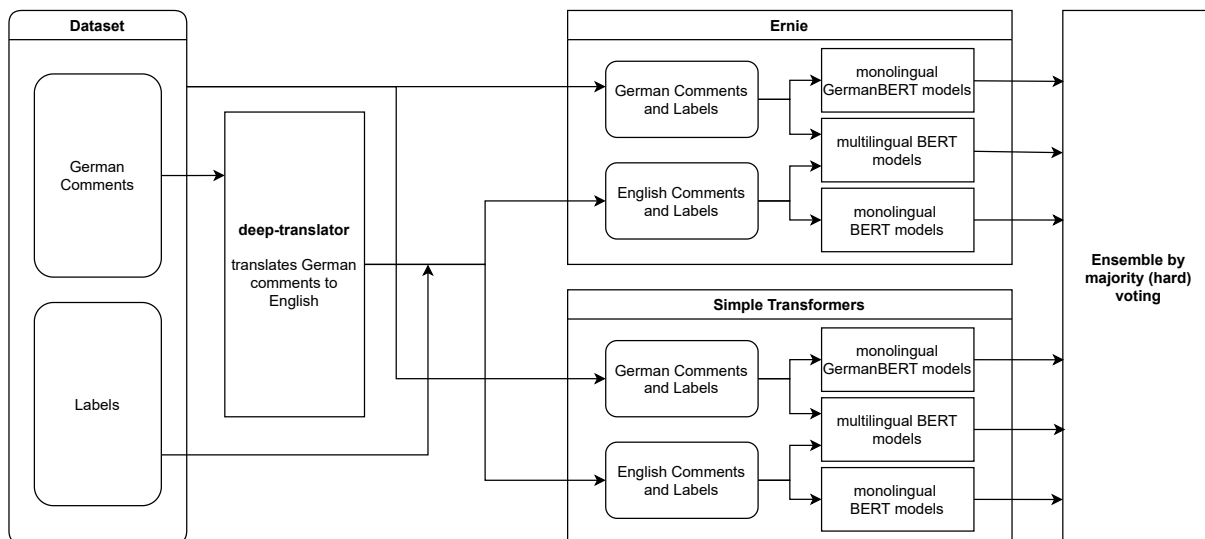


Figure 1: System architecture

Hyperparameter	Ernie	Simple Transformers
# epochs	3	3
max sequence length	128	128
learning rate	2e-5	4e-5
optimizer	Adam	AdamW

Table 2: Hyperparameter values

and evaluating the development and holdout set, the chosen models’ predictions go to the ensemble strategy, which finally predicts the test dataset by majority (hard) voting.

4.2 BERT and its variants

BERT, which stands for Bidirectional Encoder Representations from Transformers, is a language model developed by Google and is known for its state-of-the-art (SOTA) performance in several NLP tasks (Devlin et al., 2019). The Shared Task consists of German Facebook comments, so we see it fit to choose German-based and English-translation-based models. Because Facebook comments have some similarity with Twitter comments, we also decide on Twitter-based models.

There are several versions of BERT with different pre-training or fine-tuning:

- German-based BERT models
 - DBMDZ GermanBERT⁵
 - Deepset.AI GermanBERT (Chan et al., 2020)
- Multilingual BERT models

⁵<https://huggingface.co/dbmdz/bert-base-german-cased>

- mBERT Cased (Devlin et al., 2019)
- XLM-RoBERTa (Conneau et al., 2019)
- Twitter-based BERT models
 - BERTweet (Nguyen et al., 2020)
 - XLM-T (Barbieri et al., 2021)

Table 3 shows the result of each BERT model on the evaluation/holdout set and on the test dataset with its labels for subtask 1 (which was provided after the submissions had been received).

4.3 Ensembling Strategy

The Ensemble Technique is a combination of classifiers’ predictions for further classification (Opitz and Maclin, 1999). There are two popular types of ensembling: Bagging (Breiman, 1996) and Boosting (Freund and Schapire, 1999). Ensembles have been shown to be highly effective for a variety of NLP tasks, e.g., in the current top 10 of SQuAD 2.0⁶, all models are ensembles. We went for simple majority ensembling using hard voting, which classifies with the largest sum of predictions from all models.

We decided to use the three runs for the Shared Task to test different combinations of BERT models for a robust and consistent result in the test dataset. That is why we chose five models for the first run, seven models for the second run, and for the third run, nine models ensembled together. The first ensemble consists of two GermanBERT models, the English BERT_{base} model, one Twitter-based

⁶<https://rajpurkar.github.io/SQuAD-explorer/>

Classifier	Language	macro F1 _{eval}	macro F1 _{test}
1) BERT _{base} Uncased (Devlin et al., 2019)	English	.6493	.6329
2) mBERT _{base} Cased (Devlin et al., 2019)	English	.6247	.6194
3) mBERT _{base} Cased (Devlin et al., 2019)	German	.6286	.6086
4) DBMDZ GermanBERT ⁵	German	.6472	.6591
5) Deepset.AI GermanBERT (Chan et al., 2020)	German	.6481	.6608
6) BERTweet (Nguyen et al., 2020)	English	.6798	.6832
7) XLM-T (Barbieri et al., 2021)	English	.6553	.6681
8) XLM-T (Barbieri et al., 2021)	German	.6342	.6502
9) XLM-R _{base} (Conneau et al., 2019)	English	.6421	.6482
10) XLM-R _{base} (Conneau et al., 2019)	German	.3959	.3862

Table 3: BERT classifier result for subtask 1

model (BERTweet), and one multilingual model, so we have diversity for classification. For the second ensemble, one multilingual model and one Twitter-based model are added. The third ensemble has every classifier except the last one.

The results for each subtask are in Tables 4, 5, and 6, with precision, recall, and macro-averaged F1 score as the scoring metrics. The numbers in the column "Ensemble" refer to the classifier numbers from Table 3.

Run	Ensemble	P _{test}	R _{test}	macro F1 _{test}
1	1,3,4,5,6	.7047	.6588	.6810
2	1,2,3,4,5,6,8	.7183	.6635	.6898
3	1,2,3,4,5,6,7,8,9	.7168	.6529	.6833

Table 4: Ensemble result for subtask 1

Run	Ensemble	P _{test}	R _{test}	macro F1 _{test}
1	1,3,4,5,6	.7228	.6653	.6929
2	1,2,3,4,5,6,8	.7124	.6642	.6875
3	1,2,3,4,5,6,7,8,9	.7003	.6542	.6764

Table 5: Ensemble result for subtask 2

Run	Ensemble	P _{test}	R _{test}	macro F1 _{test}
1	1,3,4,5,6	.7791	.7310	.7543
2	1,2,3,4,5,6,8	.7756	.7454	.7602
3	1,2,3,4,5,6,7,8,9	.7725	.7438	.7579

Table 6: Ensemble result for subtask 3

5 Discussion

Our experiments demonstrate that BERTweet was showing better performance than every other model in every subtask, which is a surprise. We expected the monolingual GermanBERT models to perform best because of the cultural context in the integrated

German language. Multilingual BERT models perform worst but by a close margin. Because of an overfitting issue, the tenth BERT classifier XLM-R performed faultily, only recognizing negative labels and thus the low macro-averaged F1 scores. The margin of each ensemble performance in subtasks 1 and 3 is around 1%, and for subtask 2 only around 2%. We conclude that the ensembling strategy shows robustness and consistency for the choice of good classifiers in a big enough amount for each task, and it could be a legitimate approach for the overfitting problem. Because of time constraints, no cross-validation was conducted, and since the holdout set was chosen not to be released for training, there is still improvement in the training quality of the BERT models so that more experiments are needed. Each part of a system like the GPU influences the training accuracy, so an identical replication is difficult to achieve, leading to different results. That is why reproducibility is not guaranteed, even if a manual seed is set⁷. Also, the amount and the imbalance of the dataset can lead to overfitting and lower scoring.

6 Conclusion and Future Work

We presented an ensemble strategy using ten BERT classifiers, including the use of machine translation, demonstrating robustness across tasks. While ensembles perform best overall, Twitter-based models (using standard BERT hyperparameter values) with translation to English perform best in a single model setting. This observation might change if cross-validation, early stopping, hyperparameter tuning, and other optimization techniques for each model are available for future work.

⁷<https://pytorch.org/docs/stable/notes/randomness.html>

Acknowledgements

This work was supported by the project *COURAGE: A Social Media Companion Safeguarding and Educating Students* funded by the Volkswagen Foundation, grant number 95564.

References

- Phoebe Arnold. 2020. [The challenges of online fact checking](#). Technical report, Full Fact, London, UK.
- Mevan Babakar and Will Moy. 2016. [The State of Automated Factchecking](#). Technical report, Full Fact, London, UK.
- Francesco Barbieri, Luis Espinosa Anke, and José Camacho-Collados. 2021. [XLM-T: A multilingual language model toolkit for twitter](#). *CoRR*, abs/2104.12250.
- Leo Breiman. 1996. [Bagging predictors](#). *Mach. Learn.*, 24(2):123–140.
- Cambridge Consultants. 2019. [Use of AI Online in online content moderation](#).
- Branden Chan, Stefan Schweter, and Timo Möller. 2020. [German’s next language model](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6788–6796, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Unsupervised cross-lingual representation learning at scale](#). *arXiv*.
- J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Yoav Freund and Robert Schapire. 1999. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence*, 14(771-780):1612.
- Naeemul Hassan, Bill Adair, J. Hamilton, C. Li, M. Tremayne, Jun Yang, and Cong Yu. 2015. The Quest to Automate Fact-Checking. In *Proceedings of the 2015 computation+ journalism symposium*.
- Simon Kemp. 2021. [Digital 2021 april statshot report - datareportal – global digital insights](#).
- Diederik P. Kingma and Jimmy Lei Ba. 2015. [Adam: A method for stochastic optimization](#). *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, pages 1–15.
- Lev Konstantinovskiy, Oliver Price, Mevan Babakar, and Arkaitz Zubiaga. 2021. [Toward automated factchecking: Developing an annotation schema and benchmark for consistent automated claim detection](#). *Digital Threats: Research and Practice*, 2(2).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). *7th International Conference on Learning Representations, ICLR 2019*.
- Preslav Nakov, David P. A. Corney, Maram Hasanain, Firoj Alam, Tamer Elsayed, Alberto Barrón-Cedeño, Paolo Papotti, Shaden Shaar, and Giovanni Da San Martino. 2021. [Automated fact-checking for assisting human fact-checkers](#). *CoRR*, abs/2103.07769.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. [BERTweet: A pre-trained language model for English Tweets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020*, pages 9–14. Association for Computational Linguistics.
- David Opitz and Richard Maclin. 1999. [Popular Ensemble Methods: An Empirical Study](#). *Journal of Artificial Intelligence Research*, 11:169–198.
- Julian Risch and Ralf Krestel. 2020. [Top comment or flop comment? Predicting and explaining user engagement in online news discussions](#). *Proceedings of the 14th International AAAI Conference on Web and Social Media, ICWSM 2020*, pages 579–589.
- Julian Risch, Anke Stoll, Lena Wilms, and Michael Wiegand. 2021. Overview of the GermEval 2021 shared task on the identification of toxic, engaging, and fact-claiming comments. In *Proceedings of the GermEval 2021 Shared Task on the Identification of Toxic, Engaging, and Fact-Claiming Comments co-located with KONVENS*, pages 1–12.
- Guizhe Song, Degen Huang, and Zhifeng Xiao. 2021. [A study of multilingual toxic text detection approaches under imbalanced sample distribution](#). *Information (Switzerland)*, 12(5):1–16.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi 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 Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Steven Zimmerman, Chris Fox, and Udo Kruschwitz. 2019. Improving hate speech detection with deep learning ensembles. *LREC 2018 - 11th International Conference on Language Resources and Evaluation*, pages 2546–2553.