

Detecting Sexism in German Online Newspaper Comments with Open-Source Text Embeddings (Team GDA, GermEval2024 Shared Task 1: GerMS-Detect, Subtasks 1 and 2, Closed Track)

Florian Bremm¹, Patrick Gustav Blaneck¹, Tobias Bornheim², Niklas Grieger^{1,3} and Stephan Bialonski^{1,*}

¹*Department of Medical Engineering and Technomathematics*

Institute for Data-Driven Technologies, FH Aachen University of Applied Sciences, Jülich, Germany

**bialonski@fh-aachen.de*

²*ORDIX AG*

³*Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands*

Abstract

Sexism in online media comments is a pervasive challenge that often manifests subtly, complicating moderation efforts as interpretations of what constitutes sexism can vary among individuals. We study monolingual and multilingual open-source text embeddings to reliably detect sexism and misogyny in German-language online comments from an Austrian newspaper. We observed classifiers trained on text embeddings to mimic closely the individual judgements of human annotators. Our method showed robust performance in the *GermEval 2024 GerMS-Detect* Subtask 1 challenge, achieving an average macro F1 score of 0.597 (4th place, as reported on Codabench). It also accurately predicted the distribution of human annotations in *GerMS-Detect* Subtask 2, with an average Jensen-Shannon distance of 0.301 (2nd place). The computational efficiency of our approach suggests potential for scalable applications across various languages and linguistic contexts.

1 Introduction

The reliable detection of sexism and misogyny in online discussions has received increased attention in recent years (Fontanella et al., 2024). Since the events of “Gamergate” in August 2014 (Massanari, 2016), a harassment campaign targeting female journalists, research on sexism and misogyny in online platforms has gained momentum. Exposure to sexism can have tangible negative effects, for example discouraging women from participating in online discussions, as shown by an online survey conducted by an Austrian newspaper (Krenn et al., 2024). Given the often subtle and subjective nature of sexist content, moderators face significant challenges in identifying it. This highlights the need

for effective detection tools that can support moderators creating more inclusive online spaces for women.

Previous efforts to automate the detection of sexism and misogyny have typically relied on machine learning methods, often treating sexism and misogyny as forms of hate speech (Jahan and Oussalah, 2023). Numerous datasets have been created to support the training and validation of general hate speech detection models (Polletto et al., 2021; Yu et al., 2024). More recently, specialized datasets aimed specifically at sexism or misogyny detection have been released for different languages, including English (Anzovino et al., 2018), Spanish (Rodríguez-Sánchez et al., 2020), and French (Chiril et al., 2020). Alongside these developments, several competitions, such as SemEval-2019 Task 5 (Basile et al., 2019), EXIST 2022 (Rodríguez-Sánchez et al., 2022), and SemEval-2023 Task 10 (Kirk et al., 2023), have been held to promote progress in identifying sexism and misogyny. However, German-language resources for sexism detection have been particularly limited (Yu et al., 2024). The introduction of GERMS-AT (Krenn et al., 2024), a dataset of about 8000 online comments of an Austrian newspaper annotated for sexist content, has significantly improved the prospects for developing and evaluating sexism detection models in German. This dataset includes diverse annotations from multiple individuals, capturing the variability in human judgment.

In this contribution, we study the ability of open-source text embedding models, i.e., the multilingual “mE5-large”¹ (Wang et al., 2024) and the monolingual “German BERT large paraphrase

¹<https://huggingface.co/intfloat/multilingual-e5-large>

cosine² model, to reliably detect sexism and misogyny in German-language online comments (GERMS-AT). Using these text embeddings, we observed that traditional machine learning classifiers, which are fast and inexpensive to train, robustly predict the judgments of human annotators. We detail our approach and describe the models that were evaluated in the *GermEval 2024 GerMS-Detect* shared tasks and the results we obtained on out-of-sample data. The implementation details of our experiments are available online³.

2 Data and Tasks

2.1 Data

The dataset of the *GermEval 2024 GerMS-Detect* Shared Task consisted of 7984 German-language comments from the comment section of an Austrian online newspaper (Krenn et al., 2024). While all comments were in German, many comments contained Austrian dialects or slang (see Figure 1, Example 1).

Example 1:

“Des Oaschloch is eh scho berühmt, de virz’g Jungfrauen oide, notgeile Nonnen.” (ID: e1e80ff680f874d49ddf33ac846a454)

Trans.: “This asshole is already famous anyway, the forty virgins, old, horny nuns.”

No. 0-absence:	1	(A001)
No. 1-mild:	0	
No. 2-present:	1	(A007)
No. 3-strong:	5	(A002, A003, A004, A005, A012)
No. 4-extreme:	3	(A008, A009, A010)

Example 2:

“Warum wählen dann aber immer noch 36% der Frauen in Österreich die övp?” (ID: 0917bc805a3b4c3086ee7101f2740dad)

Trans.: “Then why do 36% of women in Austria still vote for the ÖVP?”

No. 0-absence:	4	(A002, A009, A010, A012)
No. 1-mild:	0	
No. 2-present:	0	
No. 3-strong:	0	
No. 4-extreme:	0	

Figure 1: Comments from the provided training dataset with annotations grouped by label (annotators shown in parentheses). The comment in *Example 1* contains an Austrian dialect and was annotated by all ten experts receiving a variety of labels. *Example 2* was only annotated by four experts and received the same label from all of them.

Each comment was annotated by at least four annotators out of a group of ten human experts

²<https://huggingface.co/deutsche-telekom/gbert-large-paraphrase-cosine>

³<https://github.com/dslaborg/germeval2024>

following specific guidelines⁴. Annotators were asked to label each comment as either not-sexist (0) or sexist (1–4). If a comment was identified as sexist, annotators assigned a label between 1 and 4, indicating the severity of the sexism or misogyny (1-mild, 2-present, 3-strong, 4-extreme). The annotations were highly subjective and varied significantly between annotators (see Figure 1). Figure 2 illustrates this variability in the annotations and highlights the imbalance in the number of comments labeled by each annotator. Additionally, Figure 2 shows, that the distribution of the five labels is highly imbalanced, with the majority of comments being labeled as not-sexist (0) by all annotators.




Figure 2: Distribution of the labels assigned by each annotator (A001–A012). Note that there are no annotations from users A006 and A011.

For the final phase of the Shared Task, the organizers provided a training dataset containing 5998 comments (75.1%) and a test dataset containing 1986 comments (24.9%), which was used for the evaluation of the final models on the competition website. The test dataset did not contain any annotations, but included the IDs of the annotators who labeled each comment. We employed two different data splits for model exploration and final training, respectively. During model exploration, we randomly split the provided training dataset into a smaller training set with 80% of the comments and a validation set containing the remaining 20%. We then created annotator-specific training sets by filtering the reduced training set for the annotations of each annotator. Furthermore, each annotator-specific training set was split into five folds for a cross-validation setup. For final training, we used the entire provided training set and, similar to the

⁴<https://ofai.github.io/GermEval2024-GerMS/guidelines.html> (Krenn et al., 2024)

model exploration phase, created annotator-specific training sets with 10% of the data reserved for early stopping.

2.2 Tasks

The Shared Task consisted of two subtasks. In *Subtask 1*, all annotations of a comment were aggregated into a single prediction target using various strategies. The goal was then to predict the aggregated label of each aggregation strategy for each comment. The following aggregation strategies were used:

- *Majority (binary prediction target)*: A comment is labeled as *not-sexist* or *sexist* depending on whether the majority of annotators labeled the comment as *not-sexist* (0) or *sexist* (1–4). If there is no majority, both labels are considered valid. In Figure 1, *Example 1* would be labeled as *sexist* and *Example 2* as *not-sexist*.
- *One (binary)*: If at least one annotator labeled a comment as *sexist* (1–4), the comment is labeled as *sexist*, otherwise as *not-sexist*. In Figure 1, *Example 1* would be labeled as *sexist* and *Example 2* as *not-sexist*.
- *All (binary)*: If all annotators labeled a comment as *sexist* (1–4), the comment is labeled as *sexist*, otherwise as *not-sexist*. In Figure 1, both examples would be labeled as *not-sexist*.
- *Majority (multi-class)*: The label of a comment is the majority label of all annotators for this comment. If there is no majority, each of the labels assigned by the annotators is considered valid. In Figure 1, *Example 1* would be labeled as *3-strong* and *Example 2* as *0-absence*.
- *Disagreement (binary)*: If at least one annotator labeled a comment as *sexist* (1–4), while at least one other annotator labeled the same comment as *not-sexist* (0), the comment is labeled as *disagreed*. Otherwise, the comment is labeled as *agreed*. In Figure 1, *Example 1* would be labeled as *disagreed* and *Example 2* as *agreed*.

In contrast to the binary and multi-class targets of *Subtask 1*, the goal of *Subtask 2* was to model the relative distribution of annotations per comment. The following two distributions were of interest:

- *Binary*: The portion of annotators labelling a comment as *not-sexist* (0) or *sexist* (1–4) respectively. In Figure 1, *Example 1* would be labeled as 10% *not-sexist* and 90% *sexist*. *Example 2* would be labeled as 100% *not-sexist*.
- *Multi-Class*: Each prediction target represents the portion of annotators labelling a comment as one of the five labels (0–4). In Figure 1, *Example 1* would be labeled as 10% *0-absence*, 0% *1-mild*, 10% *2-present*, 50% *3-strong*, and 30% *4-extreme*. *Example 2* would be labeled as 100% *0-absence*.

3 Methods and Results

Both subtasks of the *GermEval 2024 GerMS-Detect* Shared Task required knowledge of the distribution of annotations per comment. Since the test dataset contained information about the annotators that labeled the comments, we decided to train individual models for each annotator. In our approach, we combined pre-trained open-source large language models for text embeddings with simple classifiers to predict the annotations of an annotator.

3.1 Model Architecture

All comments were embedded into high-dimensional vector spaces using either (i) the “German BERT large paraphrase cosine” (GBERT-large-pc) model⁵, a version of the monolingual German-BERT model (Chan et al., 2020) that was fine-tuned for text embeddings by Deutsche Telekom or (ii) the multilingual “mE5-base”⁶ or “mE5-large”⁷ (Wang et al., 2024) models. A general overview of the models used for text embeddings is provided in Table 1.

Model	Layers	Embedding Size
mE5-base	12	768
mE5-large	24	1024
GBERT-large-pc	24	1024

Table 1: Overview of the models used for text embeddings. The models were used as is, without any further fine-tuning.

⁵<https://huggingface.co/deutsche-telekom/gbert-large-paraphrase-cosine>

⁶<https://huggingface.co/intfloat/multilingual-e5-base>

⁷<https://huggingface.co/intfloat/multilingual-e5-large>

Hyperparameter	Value/Range
Multilayer Perceptron	
hidden_layer_sizes	[64, 2048]
class_weight	oversampled
Random Forest	
n_estimators	[10, 560]
criterion	gini
max_depth	[1, 91]
class_weight	balanced
Support Vector Machine	
C	[1, 91]
kernel	rbf
class_weight	balanced

Table 2: Overview of the hyperparameter ranges used for tuning the classifiers. The hyperparameters were explored using grid search with 5-fold cross-validation. The hyperparameter “class_weight” indicates how the class imbalance was addressed.

The resulting text embeddings were then used as input features for each of the following classifiers: (i) Multilayer Perceptron (MLP), (ii) Random Forest (RFC), and (iii) Support Vector Machine (SVC).

3.2 Model Training and Evaluation

Model training consisted of a model exploration phase, where we optimized hyperparameters and selected the best-performing classifier, and the final training, where we retrained the optimal training configurations on the entire training set for submission to the competition. During training, only the classifier was updated, while the text embeddings were kept fixed. When training the Multilayer Perceptrons, we used 10% of the training data as an early stopping set to prevent overfitting.

We optimized hyperparameters separately for each annotator and classifier type using 5-fold cross-validation on each annotator’s training dataset (see Section 2.1 for details on the data split). Table 2 provides an overview of the hyperparameters of each classifier and their search ranges that were explored using grid search. To mitigate the class imbalance in the dataset (see Section 2.1), we balanced the classes for each annotator-specific dataset by either using class weights (preferred) or oversampling, depending on the respective model implementation by scikit-learn⁸ (see Table 2).

After hyperparameter tuning, we retrained mod-

els for each annotator in the best-performing configuration on the entire annotator-specific training set and evaluated them on the validation set to identify the best-performing classifier type. For evaluation, we aggregated the predictions of each annotator’s model using the respective aggregation strategies for Subtasks 1 and 2 (see Section 2.2). We then evaluated the performance of the aggregations using the Macro-F1 score for Subtask 1 and the Jensen-Shannon distance (Lin, 1991) for Subtask 2.

Once the best-performing classifier was identified, we recombined the training and validation sets and retrained the models on the entire dataset for each annotator. The final models then predicted the annotations of the test set, which were aggregated for Subtasks 1 and 2 using the same strategies as during model exploration and submitted to the competition.




Figure 3: Macro-F1 scores our models achieved when aggregating the predictions for Subtask 1 on the validation set (higher is better).

3.3 Results

Figures 3 and 4 show the performance of our models on the validation set at the end of the model exploration phase for Subtask 1 and Subtask 2, respectively. The Multilayer Perceptron and the Support Vector Machine achieved similar scores on both subtasks, with the MLP performing slightly better on Subtask 1 and the SVC on Subtask 2 for all text embedding models. In comparison, the Random Forest classifier performed worse on average on both subtasks. While model performance did not vary significantly between the different text embedding models, the mE5-large embeddings consistently outperformed the mE5-base embeddings.

⁸<https://scikit-learn.org/stable/>




Figure 4: Jensen-Shannon distances our models achieved when aggregating the predictions for Subtask 2 on the validation set (lower is better).

The embeddings from the GBERT-large-pc model performed slightly worse than the mE5-large embeddings on Subtask 1 but achieved slightly better results on Subtask 2.

Accordingly, we decided to submit variations of the Multilayer Perceptron and the Support Vector Machine classifiers with mE5-large and GBERT-large-pc embeddings to the competition. Our best-performing model for Subtask 1 was the Support Vector Machine classifier on top of mE5-large embeddings with a Macro-F1 score of 0.597, ranking 4th on the Codabench leaderboard. For Subtask 2, the Support Vector Machine classifier with GBERT-large-pc embeddings achieved the best results with an average Jensen-Shannon distance of 0.301, ranking 2nd on the leaderboard.

4 Conclusion

Our study demonstrates that support vector machines trained on open-source text embeddings can robustly predict sexism and misogyny in German-language online comments, reflecting the diversity of human judgments. While the efficiency and low cost of our training process is an advantage of our approach, we see potential for further improvements. For example, jointly fine-tuning text embeddings and classifiers on sexism datasets could improve performance, as demonstrated by the top system in the EXIST 2022 challenge, which excelled at detecting sexism in Twitter tweets (Serano, 2022). In addition, creating ensembles of fine-tuned text embedding models and classifiers could also lead to better results, similar to the strategy employed by the winning system in the GermEval

2021 challenge for identifying toxic Facebook comments (Bornheim et al., 2021). We foresee that advanced sexism detection systems can greatly assist social media moderators, paving the way for more respectful and inclusive online interactions in the future.

Acknowledgements

We are grateful to M. Reißel and V. Sander for providing us with computing resources.

References

- Maria Anzovino, Elisabetta Fersini, and Paolo Rosso. 2018. *Automatic Identification and Classification of Misogynistic Language on Twitter*, page 57–64. Springer International Publishing.
- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. *SemEval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter*. In *Proc. 13th Int. Workshop on Semantic Evaluation*, pages 54–63, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Tobias Bornheim, Niklas Grieger, and Stephan Bialonki. 2021. *FHAC at GermEval 2021: Identifying german toxic, engaging, and fact-claiming comments with ensemble learning*. *CoRR*, abs/2109.03094.
- Branden Chan, Stefan Schweter, and Timo Möller. 2020. *German’s next language model*. In *Proc. 28th Int. Conf. on Computational Linguistics, COLING 2020*, pages 6788–6796, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Patricia Chiril, Véronique Moriceau, Farah Benamara, Alda Mari, Gloria Origgi, and Marlène Coulomb-Gully. 2020. *An annotated corpus for sexism detection in french tweets*. In *Proc. 12th Language Resources and Evaluation Conf., LREC 2020*, pages 1397–1403, Marseille, France. European Language Resources Association.
- Lara Fontanella, Berta Chulvi, Elisa Ignazzi, Annalina Sarra, and Alice Tontodimamma. 2024. *How do we study misogyny in the digital age? A systematic literature review using a computational linguistic approach*. *Humanities and Social Sciences Communications*, 11(1).
- Md Saroar Jahan and Mourad Oussalah. 2023. *A systematic review of hate speech automatic detection using Natural Language Processing*. *Neurocomputing*, 546:126232.

- Hannah Kirk, Wenjie Yin, Bertie Vidgen, and Paul Röttger. 2023. [Semeval-2023 Task 10: Explainable detection of online sexism](#). In *Proc. 17th Int. Workshop on Semantic Evaluation, SemEval@ACL 2023*, pages 2193–2210. Association for Computational Linguistics.
- Brigitte Krenn, Johann Petrak, Marina Kubina, and Christian Burger. 2024. [GERMS-AT: A sexism/misogyny dataset of forum comments from an austrian online newspaper](#). In *Proc. Joint Int. Conf. on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024*, pages 7728–7739. ELRA and ICCL.
- Jianhua Lin. 1991. [Divergence measures based on the shannon entropy](#). *IEEE Trans. Inf. Theory*, 37(1):145–151.
- Adrienne Massanari. 2016. [#Gamergate and The Fap-pening: How reddit’s algorithm, governance, and culture support toxic technocultures](#). *New Media Soc.*, 19(3):329–346.
- Fabio Poletto, Valerio Basile, Manuela Sanguinetti, Cristina Bosco, and Viviana Patti. 2021. [Resources and benchmark corpora for hate speech detection: A systematic review](#). *Lang. Resour. Eval.*, 55(2):477–523.
- Francisco J. Rodríguez-Sánchez, Jorge Carrillo-de-Albornoz, and Laura Plaza. 2020. [Automatic classification of sexism in social networks: An empirical study on twitter data](#). *IEEE Access*, 8:219563–219576.
- Francisco J. Rodríguez-Sánchez, Jorge Carrillo-de-Albornoz, Laura Plaza, Adrián Mendieta-Aragón, Guillermo Marco Remón, Maryna Makeienko, María Plaza, Julio Gonzalo, Damiano Spina, and Paolo Rosso. 2022. [Overview of EXIST 2022: sexism identification in social networks](#). *Proces. del Leng. Natural*, 69:229–240.
- Alejandro Vaca Serrano. 2022. [Detecting and classifying sexism by ensembling transformers models](#). In *Proc. Iberian Languages Evaluation Forum (IberLEF 2022) co-located with the Conf. Spanish Society for Natural Language Processing (SEPLN 2022)*, volume 3202 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024. [Multilingual E5 text embeddings: A technical report](#). *CoRR*, abs/2402.05672.
- Zehui Yu, Indira Sen, Dennis Assenmacher, Mattia Samory, Leon Fröhling, Christina Dahn, Debora Nozza, and Claudia Wagner. 2024. [The unseen targets of hate - A systematic review of hateful communication datasets](#). *CoRR*, abs/2405.08562.