

# AustroTox: A Dataset for Target-Based Austrian German Offensive Language Detection

**Pia Pachinger**  
TU Wien  
pia.pachinger@tuwien.ac.at

**Janis Goldzycher**  
University of Zurich

**Anna Maria Planitzer**  
University of Vienna

**Wojciech Kusa**  
TU Wien

**Allan Hanbury**  
TU Wien

**Julia Neidhardt**  
TU Wien

## Abstract

Model interpretability in toxicity detection greatly profits from token-level annotations. However, currently such annotations are only available in English. We introduce a dataset annotated for offensive language detection sourced from a news forum, notable for its incorporation of the Austrian German dialect, comprising 4,562 user comments. In addition to binary offensiveness classification, we identify spans within each comment constituting vulgar language or representing targets of offensive statements. We evaluate fine-tuned language models as well as large language models in a zero- and few-shot fashion. The results indicate that while fine-tuned models excel in detecting linguistic peculiarities such as vulgar dialect, large language models demonstrate superior performance in detecting offensiveness in AustroTox. We publish the data and code<sup>1</sup>.

**Content warning:** This paper contains examples of offensive language to describe the annotation scheme.

## 1 Introduction

In recent years, research in the domain of content moderation has transitioned from a unidimensional to a multidimensional perspective (Cabitza et al., 2023). A one-size-fits-all approach is unable to accommodate the diverse needs of global users (Cresci et al., 2022) whose perceptions of what constitutes harmful content is contingent upon individual, contextual, and geographical factors (Bormann, 2022; Jiang et al., 2021; Kümpel and Unkel, 2023). Scholars call for less centralized and more personalized mechanisms of moderation to account for such multifaceted differences (Jhaver et al., 2023), particularly when it comes to country-specific and subsequently linguistic nuances (Jiang et al., 2021; Demus et al., 2023). Intolerant user

<sup>1</sup><https://www.pia.wien/austrotox/>,  
<https://web.ds-ifs.tuwien.ac.at/austrotox/>



Figure 1: The posts show the importance of the target of an offensive statement in determining its severity and demonstrate how non-offensive remarks can be vulgar.

comments (e.g., offensive stereotyping), in contrast to incivil user comments (e.g., vulgarity), are perceived as more offensive and as a stronger threat to democratic values and society, as well as receiving a stronger support for deletion (Kümpel and Unkel, 2023). This highlights the importance of moderation approaches that include a more nuanced understanding of online norm violations.

For determining the harmfulness of an offensive statement, its target is decisive (Bormann, 2022; Hawkins et al., 2023). Perceptions of targets of offensive comments are subject to change based on individual, contextual, cultural and intersectional factors (Hawkins et al., 2023; Jiang et al., 2021; Shahid and Vashistha, 2023) making it, therefore, crucial to effectively identify emerging targets of such statements. Figure 1 depicts examples highlighting the importance of the target of an offensive statement in determining its severity. In order to study the detection capabilities of language models in an Austrian cultural and linguistic context, we create a corpus of Austrian German comments.

Dataset	Source	#Posts	Spans	Selected Annotations
Bretschneider and Peters	FB	5,600	✗	Moderate HS, Clear HS
One Million Posts (2017)	DerStandard	11,773	✗	Inappropriate, Discriminating
Ross et al. (2017)	Twitter	469	✗	HS (scale 1-6)
GermEval 2018	Twitter	8,541	✗	Abuse, Insult, Profanity
GermEval 2019	Twitter	7,025	✗	Explicit / Implicit offense
HASOC (Mandl et al.)	Twitter, FB	4,669	✗	HS, Offensive, Profane
Assenmacher et al. (2021)	RP	85,000	✗	Insult, Profane, Threat, Racism
GermEval 2021	FB	4,188	✗	Insult, Discrimination, Vulgarity
DeTox (Demus et al., 2022)	Twitter	10,278	✗	10 target classes for HS
Multilingual HateCheck 2022	Synthetic	3,645	✗	Abuse targeted at individuals
GAHD (Goldzycher et al.)	Adversarial	10,996	✗	HS
GERMS-AT (Krenn et al.)	DerStandard	8000	✗	Sexist / Misogynous (scale 0-4)
GerDISDetect (Schütz et al.)	Media outlets	1,890	✗	11 target classes for offense
AustroTox (ours)	DerStandard	4,562	✓	Spans of Targets and Vulgarities

Table 1: Existing German datasets related to offensiveness classification and selected annotations. *HS* stands for hate speech, *FB* for Facebook, and *RP* for Rheinische Post.

Our main contributions are:

1. 4,562 user comments from a newspaper discussion forum in Austrian German annotated for offensiveness<sup>2</sup> with the article title used as context. The majority of posts is annotated by five annotators. We additionally publish the disaggregated binary offensiveness annotations.
2. Annotated spans in comments comprising targeted individuals, groups or other entities by offensive statements, and vulgarities.
3. An evaluation of fine-tuned smaller language models and large language models in a zero- and five-shot scenario.

## 2 Related Work

Research focused on identifying spans within offensive statements is primarily focused on English user comments. Examples of annotated spans in English comments are the targets of offensive statements (Zampieri et al., 2023), the spans contributing to the offensiveness label (Mathew et al., 2021; Pavlopoulos et al., 2021), and the spans comprising a violation of a moderation policy (Calabrese et al., 2022).

We list all public German datasets covering tasks related to offensiveness detection in Table 1. All

<sup>2</sup>As there are no generally accepted definitions nor distinctions for *abusiveness*, *offensiveness* and *toxicity* (Pachinger et al., 2023), we use these terms interchangeably.

German datasets containing labels related to offensiveness except for the One Million Posts and the GerMSDetect dataset focus on different varieties of German from Austrian German. AustroTox contains the same definitions for annotating vulgarities as GermEval (Risch et al., 2021), this dataset contains annotations of vulgar posts. According to their definitions, the classes *Insult* from GermEval (Wiegand et al., 2018), *Hate Speech* from DeTox and GAHD (Demus et al., 2022; Goldzycher et al., 2024), and *Offense* from HASOC (Mandl et al., 2019) can be merged into the class *Offensive* from AustroTox. This does not imply that the class *Offensive* from AustroTox can be merged into the respective classes as their definition might be more narrow. Additionally, these datasets stem from other sources than AustroTox. AustroTox is the first German dataset related to offensiveness classification containing annotated spans.

## 3 Dataset Creation

**Data source** We source AustroTox from the Austrian newspaper *DerStandard*<sup>3</sup>, a Viennese daily publication with a left-liberal stance covering domestic and international news and topics such as economy, panorama, web, sport, culture, lifestyle, science, health, education, and family. The DerStandard forum is one of the largest discourse platforms in the German-speaking world. Despite the left-liberal stance of the newspaper, this perspec-

<sup>3</sup><https://www.derstandard.at/>

tive is not reflective of the forum’s community, as DerStandard is actively working on being a low-threshold discussion platform open to everybody. As we focus on the Austrian dialect, this Austrian news media outlet’s comment sections are a suitable sample to draw from. We argue that the forum’s expansive community and the diverse range of articles and forums offered on the DerStandard website help towards minimizing bias in AustroTox. Professional moderators ensure the exclusion of hate speech which is illegal in Austria (Government, 2023) in the forum, this results in hardly any hate speech and a focus on offensive speech in the AustroTox dataset.

**Pre-filtering comments** In order to pre-filter potentially toxic comments and comments which are not considered as toxic by existing moderation technologies, we apply stratified sampling based on the toxicity score provided by the Perspective API (Jigsaw). The toxicity score is between 0 (not toxic) and 1 (severely toxic). We compute the toxicity score for 123,108 posts. Out of these posts, 873 exhibit a toxicity score between 0.9 and 1. We add these comments to the data to be annotated. Furthermore, we create the following strata defined by the toxicity score: 0-0.3, 0.3-0.5, 0.5-0.7, 0.7-0.9. Then, we randomly sample comments from each stratum. We use the following proportions for the counts of comments from the different strata: 9 : 9 : 9 : 11.

AustroTox encompasses responses to 532 articles or discussion forums on any topics covered by DerStandard. The comments were posted between November 4, 2021, and November 10, 2021. The articles and forums where the comments appear stem from a broader time period.

**Annotation campaign** We conduct the annotation with participation from master’s students specializing in Data Science and undergraduate students majoring in Linguistics, as an integral component of their academic curriculum. 30% of the annotators are registered as female through the courses registration platform, which does not necessarily mean that they self-identify as female. The majority of the annotators are Austrian and between 19 and 26 years old, annotators are required to have at least a German level of C1. The vast majority of annotators speak German as a native language. Ethical considerations pertaining to the annotation task are expounded upon in the Ethics Statement (Section 5).

The title of the article under which the comment was posted is taken into account as context when annotating the comment. While our annotation guidelines (Appendix A) include numerous examples with the intention of being prescriptive (Rottger et al., 2022), it is important to note that due to the low number of comments per annotator and the limited time allocated for training the annotators, the procedure unavoidably incorporates a subjective element.

We classify each comment as offensive or non-offensive. For non-offensive and offensive comments, we annotate spans in the text comprising vulgarities. Both, offensive and non-offensive posts may contain an unspecified number of vulgarities, as vulgar language can exist separate from offensiveness. For offensive posts, we additionally annotate spans comprising the target of the offensive statement and the type of target (Examples in Figure 1). If the target is only mentioned via a pronoun, we select the pronoun as the span comprising the target.

Adopting a definition of *vulgarity* similar to that employed by Risch et al. (2021), we define classes and spans as follows: **Offensive**: An offensive comment includes disparaging statements towards persons, groups of persons or other entities or incites to hate or violence against a person or a group of people. **Not Offensive**: A non-offensive comment does not include disparaging statements or incites to hate or violence. **Vulgarity**: Obscene, foul or boorish language that is inappropriate for civilized discourse. **Target Group**: The target of an offensive post is a group of persons or an individual insulted based on shared group characteristics. **Target Individual**: The target of an offensive post is a single person not insulted based on shared group characteristics. **Target Other**: The target of an offensive post is not a person or a group of people.

**Data aggregation** Each post is annotated by 2 to 5 annotators, the majority of posts is annotated by 5 annotators. We choose an aggregation approach that prioritizes sensitivity, where a comment requires fewer votes to be labeled as offensive compared to the number of votes needed to consider it non-offensive. A post is solely annotated as non-offensive if  $2 \leq v_n$  and  $v_o \leq \frac{v_n}{2}$ , where  $v_n$  and  $v_o$  denote the votes for the class non-offensive and offensive. The post is labelled as offensive if  $2 \leq v_o$  and  $v_n \leq \frac{2}{3} \cdot v_o$ . Posts that do not meet one of these criteria are discarded. This implies that posts

which are labelled as offensive by 3 annotators and as non-offensive by 2 annotators are labelled as offensive while posts which are labelled as offensive by 2 annotators and as non-offensive by 3 annotators are discarded. Spans comprising the different target types are annotated by majority voting of those who labelled the post as offensive. Vulgarities are annotated if  $2 \leq v$  and  $v_a \leq v + 2$ , where  $v$  denotes the number of votes for a span being a vulgarity and  $v_a$  denotes the sum of all class votes. Table 2 contains the size of AustroTox.

	Not Off	Off
Total	2,744	1,818
Not Vulgar	2,307	712
Vulgar	437	1,106
No Target		34
Target Group		869
Target Individual		572
Target Other		275
More Target Types		68

Table 2: The size of AustroTox. The fine-grained classes are determined by types of spans contained in a comment. Off stands for offensive.

**Inter Annotator Agreement** After curating 390 posts with implausible span annotations (e.g. offensive but no target), we report a Krippendorff’s Alpha of  $\alpha = 0.49$  on the binary offensiveness classification, which is comparable to related work using crowdsourcing: Sap et al. (2020) report 0.51 and Wulczyn et al. (2017) report 0.45. An  $\alpha$  of 0.5 is between random annotation ( $\alpha = 0$ ) and full agreement ( $\alpha = 1$ ). In a prescriptive annotation paradigm (Rottger et al., 2022), tentative conclusions are still acceptable with  $\alpha \geq 0.667$  (Krippendorff, 2018). While our annotation guidelines include numerous examples, it is important to note that due to the low number of comments per annotator and the limited time allocated for training the annotators, the procedure incorporates a subjective element. Care should be taken when aggregating data in cases of moderate agreement. We argue that our aggregation approach prioritizing sensitivity provides a larger decision boundary.

**Cross-validation splits** We make AustroTox available with predetermined splits for cross-validation stratified using fine-grained classes determined by the label of the post and the types of

spans it contains. The splits consist of a ratio of about 80% for training, 10% for development, and 10% for testing. Appendix B contains more details on the dataset creation.

## 4 Experiments

**Fine-tuned language models** We fine-tune and evaluate German BERT and Electra models (Chan et al., 2020) (Table 3, Appendix C). We define three tasks: Binary offensiveness classification as sequence classification, vulgarity extraction as token classification and target extraction as token classification task. For offensiveness classification, we concatenate the article title given as context and the comment as input for the models: article title: <article title> \t comment: <comment>.

**Prompted LLMs** We additionally evaluate the class and span detection capabilities of not fine-tuned LLMs. We use the following large language models for our experiments: GPT 3.5<sup>4</sup> (*gpt-3.5-turbo-1106*) (Ouyang et al., 2022), GPT 4<sup>5</sup> (*gpt-4-1106-preview*) (Achiam et al., 2023), LeoLM 7B Chat<sup>6</sup>, and Mistral<sup>7</sup> (Jiang et al., 2023). We evaluate them in a zero-shot and five-shot scenario (Table 3, Appendix C).

For the LLM evaluation, we distinguish between multitask prediction (predicting offensiveness, vulgarities and targets) and offensiveness-only classification. We create prompts that contain an offensiveness definition, article title, and the post to be classified. In the five-shot scenario, we additionally provide five titles and posts with labels that are randomly sampled from the training set for each prediction. We require the LLM to respond in JSON<sup>8</sup>. Preliminary experiments showed that only GPT-3.5 and GPT-4 were able to produce consistently valid JSON responses. We thus only evaluate these two models in the multi-task setup. For LeoLM and Mistral, we adjust the prompt, requiring them to respond with only *0* or *1*, and define the token with the higher logit as the model’s prediction.

<sup>4</sup><https://platform.openai.com/docs/models/gpt-3-5>

<sup>5</sup><https://platform.openai.com/docs/models/gpt-4-and-gpt-4-turbo>

<sup>6</sup><https://huggingface.co/LeoLM/leo-hessianai-7b-chat>

<sup>7</sup><https://huggingface.co/mistralai/Mistral-7B-v0.1>

<sup>8</sup><https://platform.openai.com/docs/guides/text-generation/json-mode>

		Offensive		Vulgar		Target		
		Post-level		Token-level		Token-level		
	Params	Binary	Macro	Binary	Macro	Micro	Macro	
BERT <sub>de</sub>	110M	.69 ± .02	.74 ± .02	.68 ± .03	.84 ± .01	.22 ± .04	.56 ± .03	
Bert-dbm		.71 ± .02	.75 ± .02	.69 ± .02	.84 ± .01	.23 ± .02	.57 ± .03	
GBERT	Base	.72 ± .02	.76 ± .01	.69 ± .04	.84 ± .02	.23 ± .02	.57 ± .02	
Gelectra		.50 ± .25	.57 ± .12	.69 ± .03	<b>.85</b> ± .01	<b>.24</b> ± .02	<b>.58</b> ± .02	
GBERT	Large	337M	.73 ± .02	.75 ± .03	<b>.71</b> ± .03	<b>.85</b> ± .01	.21 ± .12	.53 ± .16
LeoLM	0-Shot	7B	.61 ± .02	.54 ± .03	-	-	-	-
	5-Shot		.52 ± .02	.57 ± .02	-	-	-	-
Mistral	0-Shot	7.24B	.30 ± .02	.53 ± .01	-	-	-	-
	5-Shot		.55 ± .03	.64 ± .02	-	-	-	-
GPT 3.5	0-Shot	-	.68 ± .01	.64 ± .02	.40 ± .02	.69 ± .01	.17 ± .01	.52 ± .01
	5-Shot		.72 ± .01	.72 ± .02	.43 ± .02	.70 ± .01	.20 ± .01	.55 ± .01
GPT 4	0-Shot	-	.70 ± .02	.77 ± .02	.36 ± .04	.67 ± .02	.20 ± .02	.55 ± .03
	5-Shot		<b>.76</b> ± .03	<b>.81</b> ± .02	.41 ± .02	.70 ± .01	.22 ± .03	<b>.58</b> ± .03

Table 3: Mean  $F_1$  scores and standard deviations of ten-fold cross-validation of the classification tasks. We compute the Micro F1 by adding up True and False Positives and False Negatives for the three target classes.

To ensure comparability for the token-level classification tasks, we tokenize the spans generated by the GPT-models with the GBERT tokenizer.

**Evaluation outcomes** Table 3 contains the evaluation outcomes. The proprietary LLMs outperform the open-source fine-tuned models in binary offensiveness classification. We attribute the superiority of the fine-tuned models in the vulgarity detection task to the lexical nature of the vulgarity detection task. Notably, the dataset features vulgarities in Austrian dialect that are rarely encountered elsewhere. There are 437 non-offensive but vulgar comments in AustroTox. Being able to detect vulgarities can help with debiasing vulgar False Positives and vulgar False Negatives. In especially, the results suggest that marking vulgarities using fine-tuned models and then classifying the comment with marked vulgarities using GPT-4 leads to an improvement of GPT-4’s performance. Even dictionary-based detection of vulgarities might lead to an improvement of GPT-4’s performance and to more explainable results. The span annotations allow for analysis beyond comparing disagreements with binary gold labels.

The micro  $F_1$  on the targets for the four-class target classification is generally low due to a high prevalence of the negative class. The fine-tuned models slightly outperform the LLMs at detecting the targets of offensive statements.

## 5 Conclusion

We presented AustroTox, a dataset comprising user comments in Austrian German, annotated for offensiveness. We annotated spans within the comments comprising targeted individuals, groups, or other entities through offensive statements and spans comprising vulgarities. An evaluation on our dataset indicates that the smaller language models we fine-tuned and tested excel in detecting vulgar dialect, whereas the LLMs we tested demonstrate superior performance in identifying offensiveness within AustroTox.

## Acknowledgements

Pia and Anna are funded by the Vienna Science and Technology Fund (WWTF) [10.47379/ICT20015]. Janis is funded by the University of Zurich Research Priority Program project *Digital Religion(s)*. We would like to thank the students participating in the annotation campaign. Their dedication and effort have been invaluable to the success of this project. Furthermore, we would like to thank Rebekah Wegener who helped with the annotation campaign. We would like to extend our gratitude to DerStandard for sharing their data, thereby contributing significantly to the advancement of semi-automated content moderation. Lastly, the financial support by the Christian Doppler Research Association is gratefully acknowledged.

## Ethics Statement

**Annotators’ Risks** The repeated exposure of annotators to offensive content carries risks. Therefore, the annotation campaign was reviewed by the ethics committee of our institution. In the course of our work, the annotators engaged in the annotation of comments for a duration of approximately 1.5 to 3 hours. It is noteworthy that the dataset contained a higher proportion of offensive comments than the typical distribution in a user forum. The comments were sourced from a publicly accessible, moderated forum by DerStandard, ensuring that none of them could be categorized as illegal under Austrian law (Government, 2023). To mitigate potential distress, the annotators were explicitly informed that they had the option to cease annotation if they felt overwhelmed by the task without facing consequences (Appendix A.1).

**Compensation for Annotators** Participants in the annotation campaign are predominantly Master students engaged in courses focused on introductory language technology, data annotation, and natural language processing. We consider hands-on experience in annotation tasks to be highly valuable for these students, as it equips them with the necessary skills to potentially design annotation tasks in the future and to be aware of potential pitfalls and difficulties of such tasks. Moreover, we are confident that the expected workload of 1.5 to 3 hours is suitable for the participants. The annotators were informed about the publication of the data and as data annotation is a tedious task, they received a comprehensive compensation through course credits for their efforts.

**Risks of Publication of the Data** There is a potential for exploitation of our results to generate offensive online content that may elude contemporary detection systems. We believe that these risks are manageable when weighed against the improvement of the detection of offensive statements facilitated by AustroTox. We urge researchers and practitioners to uphold the privacy of the authors of posts when working with the data. And while the data is publicly available on the website of DerStandard, in order to preserve the privacy of users, we replace mentions of users and URLs with special tokens in AustroTox. DerStandard agrees to the publication of the data. Regarding copyright concerns, simple comments in online forums are usually not covered by copyright law §Section 1

UrhG (Austrian Copyright Act).

## Limitations

**Time Span and Range of Topics** The dataset comprises comments from November 4, 2021, to November 10, 2021 and therefore consists of a higher proportion of COVID-19 related topics. However, we source comments appearing in over 532 varying articles and discussion forums, ensuring diversity in topics in the dataset. The articles and forums where the comments appear stem from a broader time period. Thus, posts in our dataset refer to more events than the ones covered in between November 4, 2021, to November 10, 2021.

**Subjectivity** In the realm of human data annotation for tasks related to sentiment, a degree of subjectivity exists. Due to the small load of comments per annotator, a large pool of annotators, and limited time allocated for training the annotators, this subjective element is reflected by the dataset and is learnt by the models during the training process. We posit that the token-level annotations included in AustroTox elevate the quality of annotation by providing clearer guidance to the annotators. Furthermore, we choose an aggregation approach that prioritizes sensitivity, where a comment requires fewer votes to be labeled as offensive compared to the number of votes needed to consider it non-offensive. We posit that this method mitigates lower agreement levels. Additionally, upon publication, we publish the disaggregated annotations for binary offensiveness.

**Limitations of Experiments** Our computational experiments on the dataset are not yet exhaustive, as certain models, notably larger encoder-only Transformer models trained on German data (like GELECTRA-large), may outperform the encoder-only Transformer models examined in our study. We did not conduct hyperparameter-search which would further improve the outcomes of the evaluation. Moreover, it is essential to include more thorough testing of open-source LLMs alongside the GPT models. Our future aim is to deliver a more comprehensive evaluation on the dataset, enabling a nuanced consideration of factors influencing the performance difference of fine-tuning and zero-shot classification on larger models.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Dennis Assenmacher, Marco Niemann, Kilian Müller, Moritz Seiler, Dennis M Riehle, and Heike Trautmann. 2021. Rp-mod rp-crowd: Moderator-and crowd-annotated german news comment datasets. In *Thirty-Fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Marike Bormann. 2022. Perceptions and evaluations of incivility in public online discussions—insights from focus groups with different online actors. *Frontiers in Political Science*, 4:812145.
- Uwe Bretschneider and Ralf Peters. 2017. Detecting offensive statements towards foreigners in social media.
- Federico Cabitza, Andrea Campagner, and Valerio Basile. 2023. Toward a perspectivist turn in ground truthing for predictive computing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 6860–6868.
- Agostina Calabrese, Björn Ross, and Mirella Lapata. 2022. Explainable abuse detection as intent classification and slot filling. *Transactions of the Association for Computational Linguistics*, 10:1440–1454.
- Santiago Castro. 2017. Fast Krippendorff: Fast computation of Krippendorff’s alpha agreement measure. <https://github.com/pln-fing-udelar/fast-krippendorff>.
- 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.
- Stefano Cresci, Amaury Trujillo, and Tiziano Fagni. 2022. Personalized interventions for online moderation. In *Proceedings of the 33rd ACM Conference on Hypertext and Social Media*, pages 248–251.
- Christoph Demus, Dirk Labudde, Jonas Pitz, Nadine Probol, Mina Schütz, and Melanie Siegel. 2023. Automatische klassifikation offensiver deutscher sprache in sozialen netzwerken. In *Digitale Hate Speech: Interdisziplinäre Perspektiven auf Erkennung, Beschreibung und Regulation*, pages 65–88. Springer.
- Christoph Demus, Jonas Pitz, Mina Schütz, Nadine Probol, Melanie Siegel, and Dirk Labudde. 2022. **Detox: A comprehensive dataset for German offensive language and conversation analysis**. In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 143–153, Seattle, Washington (Hybrid). Association for Computational Linguistics.
- Janis Goldzycher, Paul Röttger, and Gerold Schneider. 2024. Improving adversarial data collection by supporting annotators: Lessons from gahd, a german hate speech dataset. *arXiv preprint arXiv:2403.19559*.
- Austrian Federal Government. 2023. **Hate postings and other punishable postings**. Accessed: 1.5.2023.
- Ian Hawkins, Jessica Roden, Miriam Attal, and Haleemah Aqel. 2023. Race and gender intertwined: why intersecting identities matter for perceptions of incivility and content moderation on social media. *Journal of Communication*, page jqad023.
- Shagun Jhaver, Alice Qian Zhang, Quan Ze Chen, Nikhila Natarajan, Ruotong Wang, and Amy X Zhang. 2023. Personalizing content moderation on social media: User perspectives on moderation choices, interface design, and labor. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW2):1–33.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Jialun Aaron Jiang, Morgan Klaus Scheuerman, Casey Fiesler, and Jed R Brubaker. 2021. Understanding international perceptions of the severity of harmful content online. *PloS one*, 16(8):e0256762.
- Google Jigsaw. **Perspective api**. Accessed: 1.4.2023.
- 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 *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 7728–7739.
- Klaus Krippendorff. 2018. *Content analysis: An introduction to its methodology*. Sage publications.
- Anna Sophie Kümpel and Julian Unkel. 2023. Differential perceptions of and reactions to incivil and intolerant user comments. *Journal of Computer-Mediated Communication*, 28(4):zmad018.
- Thomas Mandl, Sandip Modha, Prasenjit Majumder, Daksh Patel, Mohana Dave, Chintak Mandlia, and Aditya Patel. 2019. Overview of the hasoc track at fire 2019: Hate speech and offensive content identification in indo-european languages. In *Proceedings of the 11th annual meeting of the Forum for Information Retrieval Evaluation*, pages 14–17.
- 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*

- the AAAI conference on artificial intelligence*, volume 35, pages 14867–14875.
- Hiroki Nakayama. 2018. [seqeval: A python framework for sequence labeling evaluation](https://github.com/chakki-works/seqeval). Software available from <https://github.com/chakki-works/seqeval>.
- Debora Nozza and Dirk Hovy. 2023. [The state of profanity obfuscation in natural language processing scientific publications](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3897–3909, Toronto, Canada. Association for Computational Linguistics.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#).
- Pia Pachinger, Allan Hanbury, Julia Neidhardt, and Anna Planitzer. 2023. [Toward disambiguating the definitions of abusive, offensive, toxic, and uncivil comments](#). In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 107–113, Dubrovnik, Croatia. Association for Computational Linguistics.
- John Pavlopoulos, Jeffrey Sorensen, Léo Laugier, and Ion Androutsopoulos. 2021. [SemEval-2021 task 5: Toxic spans detection](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 59–69, Online. Association for Computational Linguistics.
- 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*, pages 1–12.
- Björn Ross, Michael Rist, Guillermo Carbonell, Benjamin Cabrera, Nils Kurowsky, and Michael Wojatzki. 2017. Measuring the reliability of hate speech annotations: The case of the european refugee crisis. *arXiv preprint arXiv:1701.08118*.
- Paul Röttger, Haitham Seelawi, Debora Nozza, Zeerak Talat, and Bertie Vidgen. 2022. [Multilingual Hate-Check: Functional tests for multilingual hate speech detection models](#). In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 154–169, Seattle, Washington (Hybrid). Association for Computational Linguistics.
- Paul Rottger, Bertie Vidgen, Dirk Hovy, and Janet Pierrehumbert. 2022. [Two contrasting data annotation paradigms for subjective NLP tasks](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 175–190, Seattle, United States. Association for Computational Linguistics.
- Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A. Smith, and Yejin Choi. 2020. [Social bias frames: Reasoning about social and power implications of language](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5477–5490, Online. Association for Computational Linguistics.
- Dietmar Schabus, Marcin Skowron, and Martin Trapp. 2017. One million posts: A data set of german online discussions. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1241–1244.
- Mina Schütz, Daniela PISOIU, Daria Liakhovets, Alexander Schindler, and Melanie Siegel. 2024. [GerDIS-DETECT: A German multilabel dataset for disinformation detection](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 7683–7695, Torino, Italia. ELRA and ICCL.
- Farhana Shahid and Aditya Vashistha. 2023. Decolonizing content moderation: Does uniform global community standard resemble utopian equality or western power hegemony? In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–18.
- Julia Maria Struß, Melanie Siegel, Josef Ruppenhofer, Michael Wiegand, Manfred Klenner, et al. 2019. Overview of germeval task 2, 2019 shared task on the identification of offensive language.
- Michael Wiegand, Melanie Siegel, and Josef Ruppenhofer. 2018. Overview of the germeval 2018 shared task on the identification of offensive language.
- Ellery Wulczyn, Nithum Thain, and Lucas Dixon. 2017. Ex machina: Personal attacks seen at scale. In *Proceedings of the 26th international conference on world wide web*, pages 1391–1399.
- Marcos Zampieri, Skye Morgan, Kai North, Tharindu Ranasinghe, Austin Simmmons, Paridhi Khandelwal, Sara Rosenthal, and Preslav Nakov. 2023. [Target-based offensive language identification](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 762–770, Toronto, Canada. Association for Computational Linguistics.

## A Annotation Guidelines

We depict the annotation guidelines used for the dataset creation. For the paper, we translate the example comments (except for vulgarities in dialect) with DeepL and correct the translations manually where the overall meaning of the comment is not preserved, note that we do not try to make the comment sound as if it was written by a native English speaker. We add the original comments as footnotes. Vulgar passages might not be vulgar in the



English version, nevertheless, we indicate the annotation of the original German comment. Following [Nozza and Hovy \(2023\)](#), we obfuscate vulgarities. Originally, we annotated insults and incites to hate or violence. Nevertheless, by majority voting only 60 comments were labelled as *Incite to hate or violence*, therefore, we merge the two classes into an *Offensiveness* class.

### A.1 Generalities

**Your mental health** If the comments are disturbing for you, please stop annotating and contact us. This would of course not affect your grade, we would find a solution together on how to grade you.

**Classes** There are three exclusive classes. You select a class for each comment. See Section *Classes* for details.

**Spans** There are four spans. The spans are used for tagging passages. In order to do so, mark the text you would like to tag and select one of the four tags for the spans. See Section *Spans* for details.

**Title of article** You can see the title of the article which was commented at the right side of the page. Please take it into account when classifying and tagging the comment.

**Subjectivity** It is often hard to classify a comment with one of these classes as there are many nuances of *insults* and *incites to hate and violence* (such as for example irony). If you are not sure about how to label a comment, choose the most reasonable option to you, sometimes there is no right or wrong.

**German level** If you realize that the level of German in these comments is too hard for you or there is too much dialect in them, don't hesitate to contact us, we can easily assign you the English task and you annotate the remaining comments in English.

### A.2 Class Insult

An insult includes disparaging statements towards persons or groups of persons as well as towards other entities. Insults pursue the recognisable goal of disparaging the addressee or the object of reference. **Examples:**

- *Arguments? If this group of vaccination refusers could be reached with reason and arguments, our situation would be different.*<sup>9</sup>

<sup>9</sup>Argumente?. Wenn diese Bagg\*ge von Impfverweigerern mit Vernunft und Argumenten erreichbar wäre, sähe unsere Lage anders aus.

- **Class** Insult

- **Span** Target Group: *vaccination refusers*

- **Span** Vulgarity: *group*

- *They're definitely crazy... I'm not allowed to go to the Inn, I'm not allowed to ski, I'm supposed to make up for their losses with my taxes. You're out of your mind! And no thanks, you can save yourselves the st\*pid slogans, well then just get vaccinated.*<sup>10</sup>

- **Class** Insult

- **Span** Target Group: *They*

- **Span** Vulgarity: *crazy*

- **Span** Vulgarity: *out of your mind*

- **Span** Vulgarity: *st\*pid*

- *LOL Our vaccine refusers in the football club are all pissed off because they can't go to the Christmas party now XD g\*\*ns.*<sup>11</sup>

- **Class** Insult

- **Span** Target Group: *vaccination refusers in the football club*

- **Span** Vulgarity: *g\*\*ns*.

- *Should i applaud to that sh\*t?*<sup>12</sup>

- **Class** Insult

- **Span** Target Other: *sh\*t*

- **Span** Vulgarity: *sh\*t*

### A.3 Class Incite to Hate or Violence

An incite to hate or violence against a person or a group of people. It is hard to draw the line between *insults* and *incites to hate*, as insults always somewhat incite hate. Try to decide for yourself what you actually consider hate and what is more of an insult. Insults are usually less severe. If a comment includes an insult and an incite to hate or violence, please choose the class *Incite to hate or violence*.

**Examples:**

<sup>10</sup>Die sp\*nnen doch endgültig.. Ich darf nicht ins Gasthaus, ich darf nicht Skifahren, soll deren Ausfälle mit meinen Steuern ausgleichen. Ihr h\*bt sie doch nicht alle! Und nein danke, die besch\*\*erten Sprüche na dann geh halt impfen, könnt ihr euch gleich sparen.

<sup>11</sup>Lol Unsere impfverweigerer im Fußball Verein sein voll angfressn weil's jetzt ned auf die Weihnachtsfeier gehen können XD D\*Ilos

<sup>12</sup>Soll i zu dem Sch\*\*ß auch noch applaudieren?

- *You just have to stand up to them. Yesterday you could see migrants with tools trying to cut the fence. You simply have to drive over it.* <sup>13</sup>
  - **Class** Incite to hate or violence
  - **Span** Target Group: *migrants*
- *Refugees should face the squad!*
  - **Class** Incite to hate or violence
  - **Span** Target Group: *Refugees*
- *All Austrian people are dirty!*
  - **Class** Incite to hate or violence
  - **Span** Target Group: *English people*

#### A.4 Class None

None of the above classes. **Examples:**

- *KC-GB is already a pretty weak game... Will Mahomes find his old strength at some point this season? I don't think he will this year...* <sup>14</sup>
  - **Class** None
- *they think something will happen with these rules. And the Burgenlanders look st\*pid again. That's how "motivation" works.* <sup>15</sup>
  - **Class** None
  - **Span** Vulgarity: *st\*pid*
- *yes anyway, the Ministry of Finance obviously paid for it ... wtf ^^* <sup>16</sup>
  - **Class** None
  - **Span** Vulgarity: *wtf*
- *Oh, oh, why are we sh\*t?* <sup>17</sup>
  - **Class** None
  - **Span** Vulgarity: *sh\*t*
- *She's upset about 2G but what does she suggest? Should we just let people die?* <sup>18</sup>

<sup>13</sup>Da muss man einfach hart dagegen halten. Gestern konnte man ja Migranten mit Werkzeug sehen, die den Zaun zerschneiden wollten. Da gehört einfach drübergefahren.

<sup>14</sup>KC-GB ist schon ein ziemlich schwaches Spiel...Ob Mahomes irgendwann in dieser Saison zu alter Stärke findet? Glaub heuer wird des nix mehr...

<sup>15</sup>die glauben mit diesen Regeln passiert irgendwas. Und die Burgenländer schauen wieder bl\*d aus. So funktioniert Motivation"

<sup>16</sup>ja eh hat ja offensichtlich das finanzministerium bezahlt ... wtf ^^

<sup>17</sup>Oh, oh, warum samma sch\*\*ße?

<sup>18</sup>Sie regt sich über 2G auf aber was schlägt sie denn vor? Sollen wir die Leute einfach verrecken lassen?

#### A.5 Spans

In contrast to the classes, comments can have zero, one, or multiple spans. Insults and incitements to hate or violence are targeted at an individual person, a group of persons, or something else, such as, for example, democracy (*Target Other*). Therefore, for comments classified as *Insult* or *Incite to hate or violence*, you tag at least one person or thing as *Target Individual*, *Target Group*, or *Target Other*. For the class *None*, you don't tag *Target Individual*, *Target Group*, or *Target Other*. Vulgar passages may be found in comments of all three classes.

#### A.6 Span Vulgarity

Obscene, foul or boorish language that is inappropriate or improper for civilised discourse.

- **Example** *What the F\*\*\*\*\*ck... Magnificent Interception.* <sup>19</sup>
  - **Class** None
  - **Span** Vulgarity: *What the F\*\*\*\*\*ck*

Examples of vulgar expressions: *P\*mmel, M\*st, schw\*chsinnig, Sh\*tty, Gsch\*ssene, v\*rtrottelten, D\*mming-Däumlinge, WTF, Schn\*sel, v\*rsaut, W\*ppla, K\*ffer, zum T\*\*fel, Tr\*ttel, D\*pp, D\*dl, Sch\*\*xx, d\*mn, Hosensch\*\*sser-Nerds, N\*zipack, Gsch\*ssene, Ges\*ndel, verbl\*dete, Schw\*chköpfe, sch\*\*ßen gehen, Schattenschw\*nzler, p\*ppn, P\*\*fkinesen*

#### A.7 Span Target Individual

The target of an insult or an incite to hate or violence is a single person not insulted based on shared group characteristics.

- **Example** *f\*ck you, Max Mustermann*
  - **Class** Insult
  - **Span** Target Individual: *Max Mustermann*
  - **Span** Vulgarity: *f\*ck*

#### A.8 Span Target Group

The target of an insult or an incite to hate or violence is a group of persons or an or an individual insulted based on shared group characteristics.

- **Example** *You have to treat id\*\*ts like id\*\*ts! Therefore, of course, lockdown to save the economy!* <sup>20</sup>

<sup>19</sup>What the F\*\*\*\*\*ck... Grandiose Interception.

<sup>20</sup>Mit Id\*\*ten muss man wie Id\*\*ten umgehen! Daher selbstverständlich Lockdown, um die Wirtschaft zu retten!




Figure 2: The LightTag interface

- **Class** Insult
- **Span** Target Group:  $id^{**}ts$
- **Span** Vulgarity:  $id^{**}ts$

### A.9 Span Target Other

The target of an insult or an incite is not a person or a group of people. **Examples:**

- *Please what kind of st\*pid regulation is this USRTOK*<sup>21</sup>
  - **Class** Insult
  - **Span** Target Other: *regulation*
  - **Span** Vulgarity: *st\*pid*
- *yes please f\*ck around a little longer, finally make some decisions, the actions of our government are a disaster*<sup>22</sup>
  - **Class** Insult
  - **Span** Target Other: *government*
  - **Span** Vulgarity: *f\*ck around a little longer*

## B Details on Dataset Creation

We exclusively select original comments for the dataset while excluding responses to other comments. We use the annotation tool LightTag<sup>23</sup> (Figure 2). Each annotator undertakes the task of annotating a volume ranging from 200 to 300 comments. We use Castro’s (2017) implementation for the Krippendorff’s Alpha. Mentions of users and URLs are replaced with *USRTOK* and *URLTOK*.

## C Details on Experiments

We utilize OpenAI Copilot for code implementation.

<sup>21</sup>*Bitte was ist denn das für eine d\*mme Regelung USRTOK*  
<sup>22</sup>*ja bitte sch\*\*fts noch a bisserl länger um, endlich mal Entscheidungen treffen, das Vorgehen unserer Regierung ist ein Desaster*

<sup>23</sup><https://www.lighttag.io/>

**Models** We use the following models (with licenses in parenthesis) suitable for fine-tuning for downstream tasks or for few-shot classification for our experiments:

1. BERT<sub>de</sub> cased<sup>24</sup> (MIT)
2. BERT-dbmdz<sup>25</sup> (MIT)
3. GELECTRA<sup>26</sup> (MIT)
4. GBERT-base<sup>27</sup> (MIT)
5. GBERT-large<sup>28</sup> (MIT)
6. LeoLM 7B Chat<sup>29</sup> (LLAMA 2 COMMUNITY LICENSE AGREEMENT)
7. Mistral<sup>30</sup> (Apache 2.0)
8. ChatGPT – *gpt-3.5-turbo-1106*<sup>31</sup>
9. GPT4 – *gpt-4-1106-preview*<sup>32</sup>

**Fine-Tuning** We use the Transformer-models’ implementation from the huggingface Transformers library with the default values (huggingface-hub version 0.17.3, Transformers version 4.34.0). For all models, we use a learning rate of  $5e^{-5}$ , weight decay 0.01 with 200 warm-up steps. We use a per-device train batch size of 8 examples. We train the models for a maximum of 10 epochs, with early early stopping at a patience of 3 epochs. We keep the model with the best Binary or Micro F1 score. We use four Nvidia GTX 1080 TI GPUs with 11GB RAM each to train each model. Training all offensiveness classification models took about 20 GPU hours in total, whereas vulgarity and target extraction models 16 GPU hours each.

**Prompting** Figure 3 contains the zero-shot prompt for the multitask-setup. For the Llama 2 based models, we add the Llama-style start and end spans to the prompts. To evaluate LeoLM and

<sup>24</sup><https://huggingface.co/bert-base-german-cased>

<sup>25</sup><https://huggingface.co/dbmdz/bert-base-german-cased>

<sup>26</sup><https://huggingface.co/deepset/gelectra-base>

<sup>27</sup><https://huggingface.co/deepset/gbert-base>

<sup>28</sup><https://huggingface.co/deepset/gbert-large>

<sup>29</sup><https://huggingface.co/LeoLM/leo-hessianai-7b-chat>

<sup>30</sup><https://huggingface.co/mistralai/Mistral-7B-v0.1>

<sup>31</sup><https://platform.openai.com/docs/models/gpt-3-5>

<sup>32</sup><https://platform.openai.com/docs/models/gpt-4-and-gpt-4-turbo>

### Prompt

#### System Message

Du bekommst den Titel eines Zeitungsartikels und einen Benutzerkommentar zum Artikel. Deine Aufgabe ist es, folgende Fragen zum Kommentar zu beantworten:

1. Ist der Kommentar toxisch? (1 = toxisch, 0 = nicht-toxisch)  
Definition: Ein beleidigender oder toxischer Kommentar beinhaltet verunglimpfende Äußerungen gegenüber Personen, Personengruppen oder anderen Entitäten. Er kann auch zu Hass oder Gewalt gegen einzelne Personen oder Personengruppen aufrufen.
2. Wer oder was ist das Ziel der Toxizität?  
Markiere mindestens eines der folgenden Ziele der Toxizität, falls der Kommentar toxisch ist: "Target\_Group", "Target\_Individual", oder "Target\_Other".
3. Markiere Vulgaritäten mit "Vulgarity". Vulgaritäten können in toxischen und nicht-toxischen Kommentaren vorkommen.

Antworte im JSON-Format mit den folgenden Feldern:

```
```json
{
  "Label": <0 or 1>,
  "Tags": [
    {
      "Tag": <"Target_Group", "Target_Individual", "Target_Other", or
      "Vulgarity">,
      "Token": <Span des Targets oder der Vulgarity>
    },
    ...
  ]
},..

```

#### User Message

Titel: "Impfzwang, Wurmmittel, Todesfälle: [Name]s Corona-Aussagen im Faktencheck"  
Kommentar: "Der Typ ist so ein Häus!.. Einfach unfassbar!"

### Model Output

```
{
  "Label": 1,
  "Tags": [
    {
      "Tag": "Target_Individual",
      "Token": "Der Typ"
    },
    {
      "Tag": "Vulgarity",
      "Token": "ein Häus!"
    }
  ]
}
```



Figure 3: The prompt we provide and an example response.

Mistral AI we use a cluster with eight NVIDIA GeForce RTX 3090 GPUs (24 GB RAM per GPU). We estimate two GPU hours per evaluated model.

**Evaluation** We compute the Micro F1 for the target classification using the framework of Nakayama 2018.