

# Sociocultural knowledge is needed for selection of shots in hate speech detection tasks

Antonis Maronikolakis<sup>1,2</sup> Abdullatif Köksal<sup>1,2</sup> Hinrich Schütze<sup>1,2</sup>

<sup>1</sup>Center for Information and Language Processing, LMU Munich

<sup>2</sup>Munich Center for Machine Learning

antmarakis@cis.lmu.de akoksal@cis.lmu.de

## Abstract

We introduce HATELEXICON, a lexicon of slurs and targets of hate speech for Brazil, Germany, India and Kenya, to aid model development and interpretability. First, we demonstrate how HATELEXICON can be used to interpret model predictions, showing that models developed to classify extreme speech rely heavily on target group names. Further, we propose a culturally-informed method to aid shot selection for training in low-resource settings. In few-shot learning, shot selection is of paramount importance to model performance and we need to ensure we make the most of available data. We work with HASOC German and Hindi data for training and the Multilingual HateCheck (MHC) benchmark for evaluation. We show that selecting shots based on our lexicon leads to models performing better than models trained on shots sampled randomly. Thus, when given only a few training examples, using HATELEXICON to select shots containing more sociocultural information leads to better few-shot performance. With these two use-cases we show how our HATELEXICON can be used for more effective hate speech detection.

## 1 Introduction

To curb the spread and dissemination of hate speech online, the research and industry communities have focused on the collection of hate speech data from social media and the development of models to automatically filter out harmful content.

While there have been efforts to cover multiple languages (Ousidhoum et al., 2019; Mandl et al., 2019; Ross et al., 2017; Maronikolakis et al., 2022b), most work is still conducted for English settings (Davidson et al., 2017; Founta et al., 2018; Sap et al., 2020). Concurrently, it has been shown that cross-lingual transfer capabilities of models are limited in this domain (Nozza, 2021; Ranasinghe and Zampieri, 2020), potentially due to the heavily culture-specific and subjective nature

Set	German	Hindi
Random <sub>64</sub>	0.51 <sub>3.6</sub>	0.47 <sub>2.5</sub>
Random <sub>96</sub>	0.53 <sub>4.7</sub>	0.46 <sub>3.2</sub>
Lexicon <sub>64</sub>	0.54 <sub>1.8</sub>	0.50 <sub>5.4</sub>
Lexicon <sub>96</sub>	<b>0.55<sub>1.0</sub></b>	<b>0.52<sub>1.1</sub></b>
All <sub>128</sub>	0.53 <sub>2.1</sub>	0.44 <sub>5.6</sub>

Table 1: Comparison of F1-scores for German and Hindi between randomly- and HATELEXICON-sampled training sets of sizes 64 and 96. Standard deviation in subscript.

of the data. Thus, leveraging high-resource languages to aid performance in low-resource ones is not a reliable option. Instead, methods need to be developed to better utilize the available data.

Towards efforts for more inclusive hate speech research, we are introducing HATELEXICON, a lexicon of slurs and target group denotations that can be used as an aid to model training and interpretability. Curated in collaboration with members familiar with sociopolitical balances in the examined countries (Brazil, Germany, India and Kenya), HATELEXICON aims to bring cultural knowledge to hate speech model development.

**Models often rely on keywords for predictions (Ramponi and Tonelli, 2022).** While this can be an effective tactic in developing baselines (e.g., keyword-based models), it can have undesirable effects, such as associating generally innocuous terms with extreme speech; e.g., a negative interpretation of the term ‘Muslim’. **This erroneous association between target group names and hate speech may lead to further marginalization of vulnerable groups,** misclassifying text mentioning these terms with hate speech and consequently filtering them out (Mathew et al., 2021a; Dodge et al., 2021). Further, seemingly **innocuous terms have been appropriated by extreme**

**speech peddlers and may not be picked up by models**, such as the term ‘Goldstücke’ in German (originally meaning ‘gold pieces’ and appropriated to refer to refugees in a derogatory manner). Models unable to recognize these keywords as hateful in certain contexts will lead to hate speech falling through the cracks.

This is especially salient in few-shot settings, where the wide range of targets and slurs might not be adequately captured in annotated datasets. Further, it has been shown that model performance fluctuates a lot depending on the selection of training shots (Zheng et al., 2022). Therefore, **we need a better strategy to make the most out of the available data to select shots more conducive to model performance.**

Motivated by the above problems, we propose HATELEXICON, a lexicon aiming to (i) aid model interpretability by providing ground-truth labels on common terms and (ii) improve shot selection in low-resource settings by better coverage of key terms such as targets and slurs. To create HATELEXICON, we collaborated with annotators from our examined countries, who provided a list of keywords and marked them as target groups, slurs, neutral words or any combination of these labels.

Our contributions, in short, are the following:

1. Introduce HATELEXICON, a lexicon of target group names and slurs from Brazil, Germany, India and Kenya.
2. Show how cultural information can aid in model interpretability, identifying which slurs and targets affect performance the most.
3. Show that culturally-informed sampling outperforms random sampling in few-shot hate speech detection settings (Table 1).
4. Propose a method to complement training sets by querying data using HATELEXICON terms.

## 2 Related Work

**Hate Speech Detection.** Nascent efforts to tackle hate speech focused on the curation of general-purpose, English datasets (Founta et al., 2018; Davidson et al., 2017), later expanding into more granular annotation (Guest et al., 2021; Grimmering and Klinger, 2021; Ross et al., 2017; Sap et al., 2020; Hede et al., 2021; Wiegand et al., 2021). Most work is performed on datasets of thousands of examples, allowing for straightforward

finetuning of models. In our work, we focus on low-resource settings where only a few examples are available for training and thus traditional finetuning techniques cannot be applied.

Recently, work has been conducted to cover a larger range of languages (Ousidhoum et al., 2019; Ranasinghe and Zampieri, 2020; Maronikolakis et al., 2022b; Plakidis and Rehm, 2022). In our work, we continue previous efforts into multilingual hate speech detection by proposing a lexicon of terms (pertinent to the domain hate speech) for Brazilian Portuguese, English, German, Hindi and Swahili.

Analysis has taken place both on the model and the dataset level (Mathew et al., 2021b; Wiegand et al., 2019; Madukwe et al., 2020; Kim et al., 2020; Swamy et al., 2019; Davidson et al., 2019a). Further, hate speech datasets have been examined for presence and reproduction of bias (Davidson et al., 2019b; Laugier et al., 2021; Sap et al., 2019; Maronikolakis et al., 2022a). We continue in this direction by proposing a lexicon that can aid in interpretability and model analysis efforts.

Previous work has uncovered annotator bias (Ross et al., 2017; Waseem, 2016; Posch et al., 2018; Shmueli et al., 2021; Al Kuwatly et al., 2020), with work conducted to propose frameworks of ethical data curation (Udupa et al., 2022; Jo and Gebru, 2020; Leins et al., 2020; Vidgen et al., 2019; Gebru, 2019; Mohamed et al., 2020). To mitigate bias in our work, we are directly working with community-embedded members.

Röttger et al. (2021) proposed a benchmark for unified evaluation of hate speech detection models in English, subsequently expanded into the Multilingual HateCheck (MHC) benchmark for multiple languages (Röttger et al., 2022), used in our work.

**Few-shot Learning.** Large language models exhibit zero- and few-shot capabilities (Brown et al., 2020; Wei et al., 2022; Sanh et al., 2022; Le Scao and Rush, 2021; Gao et al., 2021; Schick and Schütze, 2021b). A challenge with finetuning large language models (and few-shot learning in particular) is inconsistency: the selection of training data greatly affects performance (Zheng et al., 2022; Mosbach et al., 2021; Lu et al., 2022). Since in few-shot learning settings only a few examples are available, any noise in the data can exacerbate training issues (Köksal et al., 2022). In our work we propose a lexicon-based approach to shot selection that consistently improves performance.

Earlier work in few-shot learning focused

on prompt-based training Schick and Schütze (2021a); Fu et al. (2022); Shin et al. (2020); Logan IV et al. (2022); Zhao and Schütze (2021). Tunstall et al. (2022) introduced a prompt-free approach to learning from small datasets (SetFit). Through the use of SentenceBERT and its Siamese-network training paradigm (Reimers and Gurevych, 2019), SetFit generates pairs of training examples and learns to minimize the distance of training example representations of the same class and, conversely, to maximize the distance for examples from different classes. This process results in a model that can generate strong sentence embeddings, which can be then used to train a classification head on a task. In our work, we use SetFit to train a multilingual SentenceBERT model on German and Hindi.

### 3 Methodology

To showcase the usefulness of HATELEXICON in hate speech model development, we examine two use cases: model interpretability and few-shot model development, showcasing how HATELEXICON can be utilized to improve both processes in the hate speech domain.

#### 3.1 HATELEXICON Curation

For the curation of HATELEXICON, we employed<sup>1</sup> annotators to provide slurs, target group names and neutral words that appear often in hateful contexts online. We employed three annotators in Brazil, four in Germany, four in India and two in Kenya.

The annotators were tasked with providing terms alongside a short description. The sourcing of terms was left up to the annotators. We suggested they could use social media (e.g., searching for certain hateful hashtags or groups), but no restrictions were imposed. Instead, we relied on the sociocultural knowledge of the annotators to guide curation. We allowed for coordination between the annotators, but with no explicit instructions to actively collaborate. Terms are written in Brazilian Portuguese, English, German, Hindi or Swahili. Acceptable terms are: (i) slurs attacking the identity of a person or group, such as ethnicity, religion and sexuality. (ii) target group denotations, such as religious groups (e.g., ‘Muslim’) and marginalized communities (e.g., ‘homosexual’). (iii) neutral words that may appear often in hateful contexts

<sup>1</sup>All annotators were paid the same rate, which was above minimum wage in all countries.

Type	Brazil	Germany	India	Kenya
Neutral	30	4	3	21
Target	4	3	7	29
Slur	11	18	35	43
Neutral/Target	0	1	0	2
Neutral/Slur	0	18	1	6
Target/Slur	0	5	0	12
<i>Total</i>	<i>45</i>	<i>50</i>	<i>50</i>	<i>116</i>

Table 2: HATELEXICON statistics for terms.

Country	Text	Type	Description
Brazil	gorda	Slur	overweight women
Brazil	traveco	Slur	transsexual
Brazil	hora	Neutral	meaning ‘hour’
Germany	Flüchtling	Target	refugee
Germany	Schwuchteln	Slur	derogatory term for homosexual
Germany	Roma	Target	ethnic group
India	Bhimte	Slur	caste-ist term
India	Mullo	Slur	Muslim people
India	peaceful	Slur	Muslim people
Kenya	wakalee	Target	Kalenjin ethnic group
Kenya	nugu	Slur	generic slur
Kenya	foreskin	Slur	derogatory against uncircumcised Luo

Table 3: Example entries of HATELEXICON.

or datasets (e.g., ‘Frauenquote’, in German meaning ‘quota of/for women’). Statistics and indicative entries are shown in Table 2 and Table 3.

To evaluate the quality of our lexicon, terms submitted by one annotator were cross-checked by the other annotators of the same country. From discussions with annotator teams, it was made clear that a few terms can be assigned more than one type. For example, in German, the term ‘Schwule’ can be used by *homosexuals to describe themselves or as a slur against them*. In these instances, we allow annotation with multiple types. For example, ‘Schwule’ is annotated both as a target group denotation *and* a slur, to better capture the dual nature of the word.

#### 3.2 Interpretability

We propose the use of HATELEXICON as a tool to interpret model predictions. Popular interpretability toolkits such as LIME (Ribeiro et al., 2016) indicate which words are most associated with pre-

dictions of particular classes. In hate speech contexts, words most important for making predictions are oftentimes target group denotations or slurs. While slurs are a more obvious indicator of hateful language, target group denotations also naturally appear in hateful contexts and there is the danger of overemphasizing their association with hate speech. This correlation could lead to further marginalization of target groups, with content mentioning target group denotations being filtered out as hate speech. With HATELEXICON we can investigate keywords associated with model predictions from a more culturally-informed perspective to better verify whether the model has accrued bias against these groups.

We take as an example use case the work in XTREMESPEECH (Maronikolakis et al., 2022b), where a novel dataset of hate speech is introduced for Brazil, Germany, India and Kenya. In Maronikolakis et al. (2022b), the authors use LIME to interpret their developed mBERT models, identifying words contributing the most to predictions. In our work we operate on two levels: First, using HATELEXICON, we investigate the list of top-contributing words and show that in all examined countries, models emphasize heavily on target groups and slurs. Further, we examine the change of model representations for targets and slurs of the Kenyan and Indian subsets before and after model finetuning.

XTREMESPEECH is a hate speech dataset with social media texts collected from multiple online platforms and messaging apps. Languages covered in the dataset are Brazilian Portuguese, German, Hindi and Swahili, as well as English (either on its own or in the form of code switching with the native language).

All text in XTREMESPEECH is targeting one or more groups based on protected attributes (e.g., women or religious minorities), annotated for three levels of extremity: derogatory, exclusionary and dangerous extreme speech. Brief descriptions (as defined in (Maronikolakis et al., 2022b)) are shown below. For full definitions, we refer readers to the original paper.

1. Derogatory extreme speech: “Text that crosses the boundaries of civility within specific contexts and targets individuals/groups based on protected characteristics.”
2. Exclusionary extreme speech: “Expressions that call for or imply excluding historically

disadvantaged and vulnerable groups based on protected attributes such as national origin, gender and sexual orientation.”

3. Dangerous extreme speech: “Text that has a reasonable chance to trigger harm against target groups.”

### 3.3 Few-Shot Learning

With the general (relative) lack of non-English data in the domain of hate speech, as well as due to the difficulty of sourcing high-quality hate speech data, few-shot learning emerges as an attractive option for model development.

In few-shot learning settings, training shot selection is of great importance to model performance (Zheng et al., 2022; Köksal et al., 2022). This is especially salient in multilingual settings, where manual evaluation or prompt engineering might be challenging due to language barriers.

We propose the use of HATELEXICON to aid shot selection, allowing for more culturally-informed sampling of training examples. Instead of randomly selecting shots, we show how HATELEXICON can be used to select examples to cover a wider range of target groups and slurs in each cultural context.

We evaluate our proposed method using SetFit (Tunstall et al., 2022), training a multilingual SentenceBERT model<sup>2</sup> to discriminate between hateful and non-hateful speech.

#### 3.3.1 Data

Training data comes from HASOC (Mandl et al., 2019) and evaluation data comes from the Multilingual HateCheck benchmark (Röttger et al., 2022), on the German and Hindi subsets. HASOC is a multilingual dataset of hate speech as sourced from Twitter. We focus on the binary classification task of HASOC, where tweets are classified as either hateful or neutral. The MHC benchmark is a suite with functional tests covering a wide range of hate speech categories.

To simulate a few-shot setting, we randomly sample 128 examples (64 hateful and 64 neutral) from HASOC German and Hindi each. This forms our total training set. We sample three sets for each languages with different seeds and report averaged results. We aim to investigate whether

<sup>2</sup><https://huggingface.co/sentence-transformers/paraphrase-multilingual-mpnet-base-v2>



Set	German		Hindi	
	S	T	S	T
All <sub>128</sub>	12	10	12	9
Random <sub>32</sub>	2	0	3	3
Random <sub>64</sub>	3	4	5	4
Random <sub>96</sub>	8	8	6	4
Lexicon <sub>xx</sub>	12	10	12	9

Table 4: Distribution of slurs (S) and targets (T) in German and Hindi sets as sampled with one of the seeds.

culturally-informed shot selection (via HATELEXICON) improves performance over random shot selection. We work with three dataset sizes: 32, 64 and 96.

### 3.3.2 Shot Selection

**Sampling Method Comparison.** For random selection, we sample shots without replacement. For the lexicon-based selection, we work in two steps: (i) select all training examples that contain a slur or a target group term, (ii) further sample randomly to reach the desired training set size. In Table 4 we show the distribution of slurs and targets in each sampled set. As expected, the randomly sampled sets do contain slurs and targets, although less frequently than in the HATELEXICON-sampled sets.

**Complementing Data.** Developers tackling hate speech online might try complementing their datasets with more data to improve performance. Since data collection and annotation is expensive and challenging, especially in low-resource languages, it is imperative that the collected data is of high quality.

To simulate this setting, we are procuring more training examples using HATELEXICON, balanced between the two classes. On top of the sampled datasets as well as the entire dataset, we are further sampling from HASOC German and Hindi 16 training examples containing a target term and 16 more a slur. These 32 examples are added to the previous training sets and few-shot training are repeated. For a fair comparison, we also sample 32 examples randomly and compare performance.

With this experiment we are aiming to investigate whether we can boost performance of a given training set by collecting training data specifically containing terms from HATELEXICON. Thus, developers in need of more data can query for terms found in HATELEXICON. While in our case we are merely sampling from HASOC, an already anno-

Brazil	Germany	India	Kenya
fechar	Politiker	muslims	cows
Ucranizar	Grünen	Muslim	ruto
ucranizar	Mohammedaner	muslim	luo
safada	Juden	Muslims	wajinga
prender	Merkels	ko	kikuyu
lixo	Merkel	mullo	stupid
coisa	Regierung	Rohingyas	idiot
kkkkk	Opfer	₹	looting
Vagabundo	Islam	suvar	tangatanga
traveco	Moslems	₹	ujinga

Table 5: Top words contributing to mBERT’s predictions. Blue: target group. Red: slur. Purple: both.

tated hate speech dataset, this method could be used generally by querying for keywords on social media platforms and annotating as is practice.

## 4 Interpretability through a Cultural Lens

### 4.1 LIME Analysis

To showcase the usefulness of HATELEXICON in hate speech detection model interpretability, we analyze predictions (as reported by the authors) of XTREMESPEECH.

As part of their study, Maronikolakis et al. (2022b) conduct an interpretability analysis of mBERT predictions for a three-way classification task to identify the extremity of text (derogatory, exclusionary or dangerous). Using LIME, they identify the top-10 words contributing the most to mBERT’s predictions (shown in Table 5).

In brief, the authors conclude that target group names (such as religious groups) and slurs contribute prominently to model predictions. This exercise was performed in close collaboration with the annotators, who had to manually examine the identified top-contributing words. This process requires significant annotator effort and thus does not scale to practical settings.

With HATELEXICON, we can automate the process, significantly reducing cost and time consumption. We find that in Brazil, there are 5 slurs; in Germany, 1 slur and 4 targets; in India, 2 slurs and 5 targets and in Kenya, 1 slur and 3 targets.

It is obvious that the model relies on the presence of slurs to make decisions, since slurs are predominantly used in hateful contexts. The model, though, also relies heavily on target group denotations when making predictions. Due to the (naturally) heightened presence of target groups in hate speech training data, models might learn to associate these otherwise innocuous terms with hate

	India	Kenya
Slurs	0.25	0.19
Targets	0.25	0.22
Stop	0.22	0.22
Random	0.21	0.24

Table 6: Cosine similarity of representations between original mBERT and models finetuned on the Indian and Kenyan sets.

speech, overemphasizing their correlation with harmful content. With HATELEXICON, we are able to identify this erroneous behavior of mBERT and potentially work on mitigating this effect.

## 4.2 Change of LM Representation

To investigate the effect training on slurs and target group names has on language models, we compare mBERT’s representation of lexicon terms before and after finetuning for India and Kenya. We finetune mBERT for the three-way classification task of Maronikolakis et al. (2022b) on the Indian and Kenyan sets.<sup>3</sup> Specifically, we extract the representation of the 8th layer<sup>4</sup> for the desired tokens<sup>5</sup> and compute the cosine similarity with the corresponding representation in vanilla mBERT.

As a baseline, we compare the change of random words and stopwords from each country. Random words were sampled from the development set of Maronikolakis et al. (2022b), matching in number the HATELEXICON terms. In Table 6 we show that in Kenya, the representation of slurs changed the most after finetuning, with the representation of targets closely behind. This indicates that vanilla mBERT has not adequately learned Kenyan slurs and target groups, since their representations changes significantly after we expose the model to the terms. In India, on the other hand, the representation of slurs changed less than that of random words. Considering the low performance of the Indian models (as reported by Maronikolakis et al. (2022b)) and the fact that targets make up half the list of top-contributing words (Table 5), we hypothesize the finetuned model has not sufficiently associated slurs with extreme speech.

<sup>3</sup>Access was granted to use the Indian and Kenyan sets.

<sup>4</sup>The 8th layer has been found to contain useful representation in multilingual models (Jalili Sabet et al., 2020; Dufter and Schütze, 2020).

<sup>5</sup>When a word spans more than one token, we average the representation of each token of the word.

## 5 Few-Shot Learning

### 5.1 Setup

In our experiments, we are comparing three sets of training data: randomly-sampled (denoted with  $\text{Random}_{xx}$ ), lexicon-based sampling (denoted with  $\text{Lexicon}_{xx}$ ) and the entire training set (denoted with  $\text{All}_{xx}$ ), where  $xx$  denotes the training set size (by default equal to 128 for  $\text{All}_{xx}$ ).

Further, we denote with  $+l$  the sets complemented with 32 training examples additionally sampled using lexicon terms and we denote with  $+r$  the sets complemented with 32 training examples additionally sampled randomly.

### 5.2 Results

**Main Results.** In Figure 1, we compare macro F1-scores between HATELEXICON- and randomly-sampled training sets as well as the set containing all available training examples ( $\text{All}_{128}$ ).

In German, excluding the sets with a size of 32 which perform poorly, training sets sampled via HATELEXICON outperform the corresponding randomly-sampled sets. Both  $\text{Lexicon}_{64}$  and  $\text{Lexicon}_{96}$  outperform all randomly-sampled sets, as well as  $\text{All}_{128}$ . At the same time, with the lexicon-based sampling method, performance is more consistent across runs, especially at sizes 64 and 96 which have a small standard deviation.

The difference between lexicon- and random-based sampling is starker in the Hindi set. Lexicon-based training sets outperform all other baselines and, like in the German experiments, standard deviation is minimized when sampling with HATELEXICON, providing better stability.

It is noteworthy that HATELEXICON-based training sets regularly outperform  $\text{All}_{128}$ . We hypothesize this is due to a higher concentration of high-quality training data in  $\text{Lexicon}_{xx}$  sets. In few-shot settings, the importance of each training example is magnified and thus noise can potentially affect performance disproportionately (Zheng et al., 2022; Mosbach et al., 2021). This is in-line with other works: for example, in Schick and Schütze (2021a), results on AGNews are worse when using the largest training set. In Section 5.3, we are further investigating this phenomenon.

**Data Complementing Results.** In Table 7, we show results of our data complementing experiments. In these experiments, we are adding to the training data 32 examples sampled either randomly ( $+r$  in notation) or via HATELEXICON ( $+l$  in nota-




Figure 1: Macro F1 (areas within one standard deviation are shaded) for MHC German (a) and Hindi (b). Lexicon-based sampling (red) outperforms both random sampling (blue) and All<sub>128</sub> (orange) for set sizes 64 and 96.

tion). We see a consistent increase in performance when adding examples sampled with our proposed method. Namely, Lexicon <sub>$xx+l$</sub>  sets consistently perform better than both the original (without data complementing) baselines and the  $+r$  variants. When complementing the Random <sub>$xx$</sub>  sets, performance is slightly more inconsistent, although the  $+l$  sets still perform better than the  $+r$  variants. Further, Lexicon <sub>$xx+l$</sub>  sets have a low standard deviation, while Random <sub>$xx+l$</sub>  and Random <sub>$xx+r$</sub>  have a consistently higher standard deviation, showing that lexicon-based sampling is overall more consistent. In general, wherever we complement using HATELEXICON ( $+l$  sets), performance is better compared to both the original and the randomly-complemented ( $+r$ ) sets.

### 5.3 Ablation Study - Predicting Shots

A reason why our lexicon-based sampling method works better than random sampling may be that it retrieves less noise and fewer ambiguous examples. It has been previously shown that hard-to-learn examples (such as text that is ambiguous, misannotated or difficult to predict) do not contribute positively to model development (Swayamdipta et al., 2020). We hypothesize that our sampling method replaces a large portion of these low-quality examples with high-quality, information-rich examples.

To investigate whether our hypothesis holds true, we develop a hate speech model and *apply it on our examined training sets*. Since lexicon-sampled training sets are bound to contain more informative and unambiguous examples, they should be easier to classify correctly than randomly-sampled examples.

For this ablation study, we finetune an XLM-RoBERTa-base model (Conneau et al., 2020) separately on all the originally available German and Hindi data from HASOC (Mandl et al., 2019), excluding the 128 training examples sampled for our experiments, for a total of 2245 training examples for German and 2835 for Hindi. Then, we apply the two resulting models on our few-shot learning training sets.<sup>6</sup>

We show (Table 8) that examples sampled with our HATELEXICON are easier to classify correctly. In Hindi, the lexicon-based set is easier by 0.02 over the other sets. In German, performance on the lexicon-based set is 0.05 higher than the randomly sampled set and 0.10 higher than the entire set. Thus, we can infer that with our lexicon-based sampling method, easier examples are sourced more often than harder-to-classify ones.

Manual inspection of prediction errors of examples contained in the total training set but not in the lexicon-based set shows a high rate of low-quality text and noise (Table 9). Example 0 has been annotated as hateful even though it is just noise, containing only an account mention. Example 1 is an ambiguous (given the lack of context) example containing a sarcastic comment against a political party (die Grünen / Greens), classified by the model as hateful. Example 2 is also ambiguous, containing sarcasm against a right-wing party in Germany (AfD). All ambiguous examples (1-3) are short tweets that mention political entities.<sup>7</sup> In a politically charged environment, these short texts

<sup>6</sup>We only predict shots from a randomly-chosen training set for each language instead of all three used previously.

<sup>7</sup>While political entities are an integral part of society, they are not target groups and were not added to our lexicon.

Set	F1	$\Delta$	Set	F1	$\Delta$
Random <sub>32+l</sub>	0.54 <sub>3.2</sub>	+0.05	Random <sub>32+l</sub>	0.47 <sub>3.3</sub>	+0.05
Random <sub>32+r</sub>	0.51 <sub>3.6</sub>	+0.02	Random <sub>32+r</sub>	0.44 <sub>4.2</sub>	+0.02
Random <sub>64+l</sub>	0.54 <sub>2.1</sub>	+0.03	Random <sub>64+l</sub>	0.50 <sub>1.8</sub>	+0.03
Random <sub>64+r</sub>	0.54 <sub>1.9</sub>	+0.03	Random <sub>64+r</sub>	0.48 <sub>2.4</sub>	+0.03
Random <sub>96+l</sub>	0.56 <sub>1.2</sub>	+0.03	Random <sub>96+l</sub>	0.51 <sub>1.1</sub>	+0.03
Random <sub>96+r</sub>	0.51 <sub>0.9</sub>	-0.02	Random <sub>96+r</sub>	0.48 <sub>2.9</sub>	-0.02
Lexicon <sub>32+l</sub>	0.56 <sub>1.2</sub>	+0.14	Lexicon <sub>32+l</sub>	0.50 <sub>0.9</sub>	+0.14
Lexicon <sub>32+r</sub>	0.51 <sub>3.0</sub>	+0.09	Lexicon <sub>32+r</sub>	0.48 <sub>2.0</sub>	+0.09
Lexicon <sub>64+l</sub>	0.56 <sub>0.9</sub>	+0.02	Lexicon <sub>64+l</sub>	0.52 <sub>0.3</sub>	+0.02
Lexicon <sub>64+r</sub>	0.55 <sub>1.8</sub>	+0.01	Lexicon <sub>64+r</sub>	0.51 <sub>0.9</sub>	+0.01
Lexicon <sub>96+l</sub>	0.59 <sub>0.4</sub>	+0.04	Lexicon <sub>96+l</sub>	0.55 <sub>0.3</sub>	+0.04
Lexicon <sub>96+r</sub>	0.56 <sub>0.9</sub>	+0.01	Lexicon <sub>96+r</sub>	0.51 <sub>0.6</sub>	+0.01
All <sub>128+l</sub>	0.57 <sub>1.1</sub>	+0.04	All <sub>128+l</sub>	0.49 <sub>1.3</sub>	+0.04
All <sub>128+r</sub>	0.54 <sub>2.2</sub>	+0.01	All <sub>128+r</sub>	0.49 <sub>1.0</sub>	+0.01

(a) German

(b) Hindi

Table 7: Macro F1 (standard deviation as subscript) and difference with the non-complemented baseline ( $\Delta$ ), for MHC German (a) and Hindi (b).

	Germany	India
Lexicon	0.61	0.55
Random	0.56	0.53
All	0.51	0.53

Table 8: F1-score of classifying training shots.

ID	Text	Type
0	@Hartes_Geld	Noise
1	Ja so tierlieb sind die grünen	Ambiguous
2	@SaschaUlbrich @Mundauf- machen @AfD super, gut gemacht! auf jeden Fall "retweeten"!	Ambiguous
3	Wer soll jetzt die SPD führen?	Low-content

Table 9: Manual inspection of model prediction errors.

do not provide enough context for the model to adequately learn whether the example is hateful or not. Therefore, adding these examples in our training set is not beneficial.

## 6 Conclusion

In our work, we curate HATELEXICON, a lexicon of slurs and targets of hate speech for the countries of Brazil, Germany, India and Kenya, with the goal of improving model development.

With our lexicon, we show how models rely on slurs and target group denotations when mak-

ing predictions in hate speech tasks. The over-reliance on target group names may lead to further marginalization of targets of hate speech, with models flagging as hateful innocuous text containing these terms. With HATELEXICON, this erroneous behavior is unveiled and researchers can focus on mitigating this bias.

We also demonstrate how HATELEXICON can be used for few-shot learning. We evaluate on the German and Hindi subsets of the Multilingual HateCheck benchmark (Röttger et al., 2022) and show that selecting training shots with a culturally-informed process (e.g., our lexicon of slurs and targets) can aid the development of hate speech classifiers. Namely, training sets sampled using HATELEXICON perform better than training sets sampled at random.

More abstractly, we provide evidence that curating sociocultural knowledge bases (e.g., lexicons) is pivotal in developing hate speech detection models. Sociocultural information is vital in contextualizing hate speech, and without it we risk developing models detached from the reality and experiences of the most vulnerable. Thus, we advocate for a greater focus on bridging the knowledge gap between researchers and affected communities for the development of models better geared towards protecting target groups.

**Acknowledgments.** This work was funded by the European Research Council (grant #740516).



## 7 Ethical Considerations and Limitations

### 7.1 Ethics Statement

In our work we are dealing with sensitive content in the form of hate speech against marginalized communities. We are not advocating for hate speech, but instead propose methods to aid in filtering out harmful content from online spheres and analyzing detection models with our proposed lexicon.

The lexicon was developed in collaboration with annotators familiar with sociocultural balances in their countries and communities, with the goal of creating a dictionary of terms useful for hate speech model development. A potential concern with a dictionary of hateful terms is that the terms will be publicized and could be subsequently used by hate speech peddlers to cause further harm. Since these terms were recorded specifically because they are already used extensively, the risk of additional harm from publicizing these terms is minimal. Moreover, in HATELEXICON we are collecting denotations of target groups. These may be based on ethnicity, religion, sexuality or other protected attributes. A potential concern is that we will be exposing the mentioned target groups. We argue that better understanding the harms faced by these communities outweighs the negatives and will provide more net-positive in the long term, while at the same time these groups were recorded due to the increased quantity of hateful content they receive.

### 7.2 Limitations

The lists of slurs and target groups in HATELEXICON are not exhaustive. While we took care to expand HATELEXICON as thoroughly as possible, we are limited by time and resources and could only cover a partial set of terms used online in relation to hate speech in the examined countries.

Further, the list of countries chosen is small: Brazil, Germany, India and Kenya. Ideally we would have included more countries and languages. More work needs to be done to expand this list and provide more coverage.

## References

Hala Al Kuwatly, Maximilian Wich, and Georg Groh. 2020. [Identifying and measuring annotator bias based on annotators' demographic characteristics](#). In *Proceedings of the Fourth Workshop on Online*

*Abuse and Harms*, pages 184–190, Online. Association for Computational Linguistics.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.

Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. 2019a. [Racial bias in hate speech and abusive language detection datasets](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 25–35, Florence, Italy. Association for Computational Linguistics.

Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. 2019b. [Racial bias in hate speech and abusive language detection datasets](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 25–35, Florence, Italy. Association for Computational Linguistics.

Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. [Automated hate speech detection and the problem of offensive language](#). In *International AAAI Conference on Web and Social Media*.

Bruno Heridet Delapouite. Accessed 10/11/2021. <https://game-icons.net/>.

Jesse Dodge, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, Margaret Mitchell, and Matt Gardner. 2021. [Documenting large webtext corpora: A case study on the colossal clean crawled corpus](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1286–1305, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Philipp Dufter and Hinrich Schütze. 2020. [Identifying elements essential for BERT's multilinguality](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4423–4437, Online. Association for Computational Linguistics.

- Antigoni-Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large scale crowdsourcing and characterization of twitter abusive behavior. In *11th International Conference on Web and Social Media, ICWSM 2018*. AAAI Press.
- Jinlan Fu, See-Kiong Ng, and Pengfei Liu. 2022. [Polyglot prompt: Multilingual multitask prompttraining](#).
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.
- Timnit Gebru. 2019. [Oxford handbook on ai ethics book chapter on race and gender](#).
- Lara Grimminger and Roman Klinger. 2021. [Hate towards the political opponent: A Twitter corpus study of the 2020 US elections on the basis of offensive speech and stance detection](#). In *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 171–180, Online. Association for Computational Linguistics.
- Ella Guest, Bertie Vidgen, Alexandros Mittos, Nishanth Sastry, Gareth Tyson, and Helen Margetts. 2021. [An expert annotated dataset for the detection of online misogyny](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1336–1350, Online. Association for Computational Linguistics.
- Anushree Hede, Oshin Agarwal, Linda Lu, Diana C. Mutz, and Ani Nenkova. 2021. [From toxicity in online comments to incivility in American news: Proceed with caution](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2620–2630, Online. Association for Computational Linguistics.
- Masoud Jalili Sabet, Philipp Dufter, François Yvon, and Hinrich Schütze. 2020. [SimAlign: High quality word alignments without parallel training data using static and contextualized embeddings](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1627–1643, Online. Association for Computational Linguistics.
- Eun Seo Jo and Timnit Gebru. 2020. [Lessons from archives: Strategies for collecting sociocultural data in machine learning](#). In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT\* '20*, page 306–316, New York, NY, USA. Association for Computing Machinery.
- Jae-Yeon Kim, Carlos Ortiz, Sarah Nam, Sarah Santiago, and Vivek Datta. 2020. [Intersectional bias in hate speech and abusive language datasets](#). *CoRR*, abs/2005.05921.
- Abdullatif Köksal, Timo Schick, and Hinrich Schütze. 2022. [Meal: Stable and active learning for few-shot prompting](#).
- Léo Laugier, John Pavlopoulos, Jeffrey Sorensen, and Lucas Dixon. 2021. [Civil rephrases of toxic texts with self-supervised transformers](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1442–1461, Online. Association for Computational Linguistics.
- Teven Le Scao and Alexander Rush. 2021. [How many data points is a prompt worth?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2627–2636, Online. Association for Computational Linguistics.
- Kobi Leins, Jey Han Lau, and Timothy Baldwin. 2020. [Give me convenience and give her death: Who should decide what uses of NLP are appropriate, and on what basis?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2908–2913, Online. Association for Computational Linguistics.
- Robert Logan IV, Ivana Balazevic, Eric Wallace, Fabio Petroni, Sameer Singh, and Sebastian Riedel. 2022. [Cutting down on prompts and parameters: Simple few-shot learning with language models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2824–2835, Dublin, Ireland. Association for Computational Linguistics.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022. [Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, Dublin, Ireland. Association for Computational Linguistics.
- Kosisochukwu Madukwe, Xiaoying Gao, and Bing Xue. 2020. [In data we trust: A critical analysis of hate speech detection datasets](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 150–161, Online. Association for Computational Linguistics.
- 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 forum for information retrieval evaluation*, pages 14–17.
- Antonis Maronikolakis, Philip Baader, and Hinrich Schütze. 2022a. [Analyzing hate speech data along](#)

- racial, gender and intersectional axes. In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 1–7, Seattle, Washington. Association for Computational Linguistics.
- Antonis Maronikolakis, Axel Wisioerek, Leah Nann, Haris Jabbar, Sahana Udupa, and Hinrich Schuetze. 2022b. [Listening to affected communities to define extreme speech: Dataset and experiments](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1089–1104, Dublin, Ireland. Association for Computational Linguistics.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2021a. [Hatexplain: A benchmark dataset for explainable hate speech detection](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17):14867–14875.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2021b. [Hatexplain: A benchmark dataset for explainable hate speech detection](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17):14867–14875.
- Shakir Mohamed, Marie-Therese Png, and William Isaac. 2020. [Decolonial ai: Decolonial theory as sociotechnical foresight in artificial intelligence](#). *Philosophy and Technology*, 33(4):659–684.
- Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. 2021. [On the stability of fine-tuning {bert}: Misconceptions, explanations, and strong baselines](#). In *International Conference on Learning Representations*.
- Debora Nozza. 2021. [Exposing the limits of zero-shot cross-lingual hate speech detection](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 907–914, Online. Association for Computational Linguistics.
- Nedjma Ousidhoum, Zizheng Lin, Hongming Zhang, Yangqiu Song, and Dit-Yan Yeung. 2019. [Multilingual and multi-aspect hate speech analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4675–4684, Hong Kong, China. Association for Computational Linguistics.
- Melina Plakidis and Georg Rehm. 2022. [A dataset of offensive German language tweets annotated for speech acts](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4799–4807, Marseille, France. European Language Resources Association.
- Lisa Posch, Arnim Bleier, Fabian Flöck, and Markus Strohmaier. 2018. [Characterizing the global crowd workforce: A cross-country comparison of crowd-worker demographics](#).
- Alan Ramponi and Sara Tonelli. 2022. [Features or spurious artifacts? data-centric baselines for fair and robust hate speech detection](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3027–3040, Seattle, United States. Association for Computational Linguistics.
- Tharindu Ranasinghe and Marcos Zampieri. 2020. [Multilingual offensive language identification with cross-lingual embeddings](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5838–5844, Online. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#).
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. ["why should i trust you?": Explaining the predictions of any classifier](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, page 1135–1144, New York, NY, USA. Association for Computing Machinery.
- Björn Ross, Michael Rist, Guillermo Carbonell, Benjamin Cabrera, Nils Kurovsky, and Michael Wotatzki. 2017. [Measuring the reliability of hate speech annotations: The case of the european refugee crisis](#). *CoRR*, abs/1701.08118.
- Paul Röttger, Haitham Seelawi, Debora Nozza, Zeerak Talat, and Bertie Vidgen. 2022. [Multilingual HateCheck: 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 Röttger, Bertie Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet Pierrehumbert. 2021. [HateCheck: Functional tests for hate speech detection models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 41–58, Online. Association for Computational Linguistics.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries,



- Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M Rush. 2022. [Multitask prompted training enables zero-shot task generalization](#). In *International Conference on Learning Representations*.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. [The risk of racial bias in hate speech detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence, Italy. 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.
- Timo Schick and Hinrich Schütze. 2021a. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021b. [It’s not just size that matters: Small language models are also few-shot learners](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352, Online. Association for Computational Linguistics.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. [AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235, Online. Association for Computational Linguistics.
- Boaz Shmueli, Jan Fell, Soumya Ray, and Lun-Wei Ku. 2021. [Beyond fair pay: Ethical implications of NLP crowdsourcing](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3758–3769, Online. Association for Computational Linguistics.
- Steve Durairaj Swamy, Anupam Jamatia, and Björn Gambäck. 2019. [Studying generalisability across abusive language detection datasets](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 940–950, Hong Kong, China. Association for Computational Linguistics.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. [Dataset cartography: Mapping and diagnosing datasets with training dynamics](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293, Online. Association for Computational Linguistics.
- Lewis Tunstall, Nils Reimers, Unso Eun Seo Jo, Luke Bates, Daniel Korat, Moshe Wasserblat, and Oren Pereg. 2022. [Efficient few-shot learning without prompts](#).
- Sahana Udupa, Antonis Maronikolakis, Hinrich Schütze, and Axel Wisioerek. 2022. [Ethical scaling for content moderation: Extreme speech and the \(in\)significance of artificial intelligence](#).
- Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019. [Challenges and frontiers in abusive content detection](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 80–93, Florence, Italy. Association for Computational Linguistics.
- Zeerak Waseem. 2016. [Are you a racist or am I seeing things? annotator influence on hate speech detection on Twitter](#). In *Proceedings of the First Workshop on NLP and Computational Social Science*, pages 138–142, Austin, Texas. Association for Computational Linguistics.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent abilities of large language models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Michael Wiegand, Josef Ruppenhofer, and Elisabeth Eder. 2021. [Implicitly abusive language – what does it actually look like and why are we not getting there?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 576–587, Online. Association for Computational Linguistics.
- Michael Wiegand, Josef Ruppenhofer, and Thomas Kleinbauer. 2019. [Detection of Abusive Language: the Problem of Biased Datasets](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 602–608, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mengjie Zhao and Hinrich Schütze. 2021. [Discrete and soft prompting for multilingual models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8547–8555, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yanan Zheng, Jing Zhou, Yujie Qian, Ming Ding, Chonghua Liao, Li Jian, Ruslan Salakhutdinov,



Jie Tang, Sebastian Ruder, and Zhilin Yang. 2022. [FewNLU: Benchmarking state-of-the-art methods for few-shot natural language understanding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 501–516, Dublin, Ireland. Association for Computational Linguistics.