

LAHM : Large Annotated Dataset for Multi-Domain and Multilingual Hate Speech Identification

Ankit Yadav, Shubham Chandel, Sushant Chatufale and Anil Bandhakavi

Logically.ai, Brookfoot Mills, Brookfoot Industrial Estate, Brighouse, HD6 2RW, United Kingdom

Abstract

Current research on hate speech analysis is typically oriented towards monolingual and single classification tasks. In this paper, we present a new multilingual hate speech analysis dataset for English, Hindi, Arabic, French, German and Spanish languages for multiple domains across hate speech - Abuse, Racism, Sexism, Religious Hate and Extremism. To the best of our knowledge, this paper is the first to address the problem of identifying various types of hate speech in these five wide domains in these six languages. In this work, we describe how we created the dataset, created annotations at high level and low level for different domains and how we use it to test the current state-of-the-art multilingual and multitask learning approaches. We evaluate our dataset in various monolingual, cross-lingual and machine translation classification settings and compare it against open source English datasets that we aggregated and merged for this task. Then we discuss how this approach can be used to create large scale hate-speech datasets and how to leverage our annotations in order to improve hate speech detection and classification in general.

Keywords

hate speech, multilingual, multi-domain, cross-lingual, racism, religious hate, sexism, abuse, extremism, few shot learning, zero shot learning.

1. Introduction

Abusive language is an important and relevant issue in social media platforms such as Twitter. Social media is often exploited to propagate toxic content such as hate speech or other forms of abusive language. The amount of user-generated content produced every minute is very large, and manually monitoring abusive behavior in Twitter is not feasible and impractical. Twitter has made efforts to eliminate abusive content from their platform by providing clear policies on hateful conduct, user reporting and using moderators to filter content. Still, these manual efforts are not scalable enough and are not long term.

Several studies from the Natural Language Processing (NLP) field have been done to tackle the problem of hate speech detection in social media. Most studies proposed a supervised approach to detect abusive content automatically using various models ranging from traditional machine learning approaches to deep learning based approaches. However, the majority of work focused only on a single language, i.e., English, and a single abusive domain phenomenon, e.g., hate


De-Factify: Workshop on Multimodal Fact Checking and Hate Speech Detection, co-located with AAAI 2022. 2022 Vancouver, Canada

✉ ankit.yadav@logically.co.uk (A. Yadav); Shubham1_c@logically.co.uk (S. Chandel); sushant.c@logically.co.uk (S. Chatufale); anil@logically.co.uk (A. Bandhakavi)

🌐 <https://www.logically.ai/team/leadership/anil-bandhakavi> (A. Bandhakavi)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

speech, sexism, racism, religious hate and so on, rather than multiple languages and multiple domains. Twitter supports content in 34 languages and user can use any one of them to express views. Thus the problem of tackling hateful content in real time in multiple languages becomes a challenge. We need robust models for hateful content detection across multiple languages and multiple domains.

In this paper we try to tackle two prominent challenges in hate speech detection-

1. Build a **multilingual dataset** for hate speech detection across 6 languages -: *English, Hindi, French, Arabic, German, Spanish*.
2. Build a **multi-domain dataset** that covers these hate speech domains -: *Abusive, Racism, Sexism, Religious Hate and Extremism*

We define the different domains as follows:-

1. RACISM: Discrimination based on race, ethnicity, caste, nationality, culture, skin colour, hair texture, physical aspects.
2. SEXISM: Discrimination based on gender/sexual orientation.
3. RELIGIOUS HATE: Religious discrimination treating a person or group differently because of the particular faith/belief which they hold about a religion.
4. ABUSE: Speech that causes or likely to cause distress, disrespect or mental pain, especially from vulgar and profane comments.
5. EXTREMISM: Speech that cause or is likely to cause, harm to individuals, communities or wider society, and where any political, civil issues can lead to extremist behaviour through violence.

2. Related Work

There have been several studies on abusive language detection [1] [2], offensive language, hate speech identification[3], toxicity [4], hatefulness [5], aggression [6], attack [7], racism, sexism [8], obscenity, threats, and insults.

Along with that, there are several shared tasks that have focused on abusive language and hate speech detection such as HASOC-2019 [9], TRAC shared task on aggression identification [6], HatEval [10] and GermEval-2018 [11] which focused on offensive language identification in German tweets.

Waseem[8] proposed the following list to identify hate speech. Their criteria are partially derived by negating the privileges observed in McIntosh (2003), where they occur as ways to highlight importance, ensure an audience, and ensure safety for white people, and partially derived from applying common sense. A tweet is categorized as offensive if it:

1. uses a sexist or racial slur.
2. attacks a minority.
3. seeks to silence a minority.
4. criticizes a minority (without a well founded argument).
5. promotes, but does not directly use, hate speech or violent crime.
6. criticizes a minority and uses a straw man argument.

7. blatantly misrepresents truth or seeks to distort views on a minority with unfounded claims.
8. shows support of problematic hash tags - “#BanIslam”, “#whoriental”, “#whitegenocide”
9. negatively stereotypes a minority.
10. defends xenophobia or sexism.
11. contains a screen name that is offensive, as per the previous criteria, the tweet is ambiguous (at best), and the tweet is on a topic that satisfies any of the above criteria.

Wide variety of machine learning models have been used to deal with multi-domain hate speech detection task. Some studies use traditional machine learning approaches such as, logistic regression [12], support vector machine [13] [14], linear support vector machine classifiers (LSVC) [15] [16] [17]. They used it for better explainability, along with several deep learning based models, including convolutions neural networks [18] [19], LSTM [20] [18] [17] [16] [21], bidirectional LSTM [22]. Most recent works focus on transfer learning and novel architectures involving Transformers based models such as Bidirectional Encoder Representations from Transformers (BERT) [23] [24] [25] [26] [27] and its variants like RoBERTa [25] in the cross-domain abusive language detection task.

For cross-lingual abusive language detection, most studies utilized transformers based models. Some traditional models used such as logistic regression [28] [10] [29], linear support vector machines [16] [17], SVM [30], LSTM [31] [29] and Bi-LSTMs [31] have also been used. Recent work focused on several transformer based architectures such as multilingual BERT [32] [28] [25] [17] [33] [34] [29], RoBERTa [35], XLM [31] [34] and XLM-RoBERTa [35] [25] [36]. Transformers based models with multilingual language representations can easily deal with language shift in zero-shot cross-lingual task.

3. LAHM Dataset

In this section, we describe the characteristics that we want to include in our dataset, our approach to collect different types of hate speech while covering all major hate speech domains and how to annotate data at large scale. We also give detailed statistics and analysis for the collected data.

3.1. Essentials of LAHM dataset

Considering no such dataset is available at present that covers these five domains of hate speech in the six languages, our aim is to create a reliable multilingual and multifaceted hate speech dataset.

Multilingual data: Our dataset is created as multilingual resource to facilitate cross-lingual research. It contains hate speech in English, Hindi, French, Arabic, German and Spanish languages.

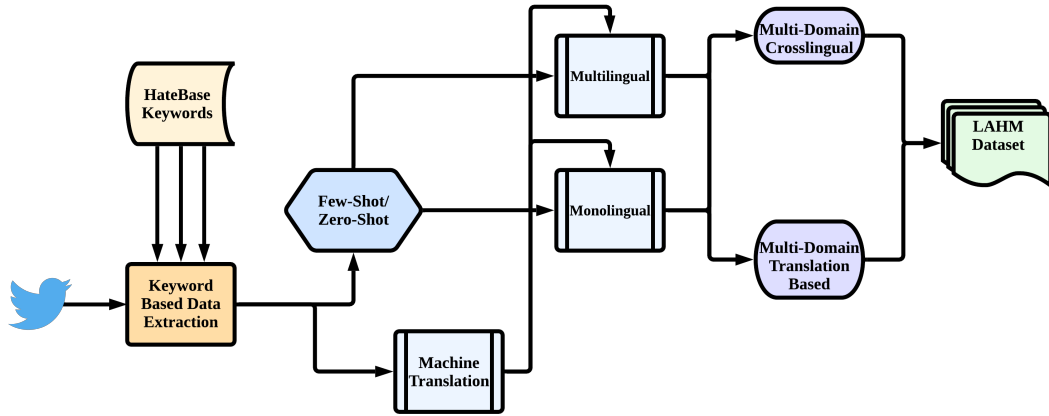


Figure 1: Overall Annotation Pipeline Architecture

Multi-Domain data: Our dataset consists of fine-grained labels for each hate sample per language. These aspects cover majority of hate domains such as racism, sexism, abuse, religious hate and extremism.

3.2. Dataset Collection

3.2.1. Language and Domain Specific Keywords Collection

Considering the cultural differences in the main regions where English, Hindi, French, Arabic, Spanish and German are spoken, we start by looking for the hateful keywords that are native to these languages. We use HateBase vocabulary dataset, which is a valuable lexicon for creating hateful dataset from public forums, as well as Hatebase’s sightings dataset, which is useful for trending analysis of keywords.

We targeted 6 languages. Additionally we tried some indic languages mainly Marathi and Bangla but hatebase has negligible coverage for these languages in terms of keywords. Additionally, we extracted the following:

1. Targeted groups for the keywords
2. Offensive levels of keywords - *Extremely, Highly and Mildly offensive.*
3. Recent sighting counts

Target groups help us categorise keywords in low level classes of ethnicity & nationality, religious hate, gender & sexual orientation and general abuse. Details of keywords per domain and language are provided in Table 1 in dataset statistics. Total keywords collected were **1022** . To extract data for extremism, we used a set of keywords related to extremism and terrorism (including terrorist organisation names) to retrieve news articles from BBC Monitoring¹.

¹<https://monitoring.bbc.co.uk/>

1. We further extracted comments from those articles related to extremism and terrorism. We also extracted tweets and comments which were extremist in nature from the counterextremism database².
2. The comments extracted were said by members and suspects of terrorist and extremist organisations on Facebook, Twitter and YouTube.

3.2.2. Large Scale Multilingual Dataset of Tweets

We started our dataset collection by using the keywords built per language. We utilize twint³ API to collect 1000-2000 tweets per keyword as Twitter official API can only be utilized for limited number of requests. We searched through the API for last one year of data. We also add additional keywords from MLMA [1] for Arabic and French.

The collected raw data contains cross-lingual tweets and therefore language detection becomes a part of our process. For each keyword in each language we consider tweets only in that language and drop the rest. This helps us in training monolingual hate speech detection without any need to worry for code switching in languages.

We substituted all usernames with @USER and urls with @URL and cleaned any unnecessary symbols from the tweet. We also discarded short tweets with less than 4 words. In total we collected **497660** tweets. Details of tweets per language is given in section 3.3 *Dataset Statistics*.

For extremism class we carefully handcrafted a total of **88** keywords and collected the data from BBC monitoring. Additional data was collected from other websites for extremism.

3.3. Dataset Statistics

Table 1 gives us the distribution of keywords across different languages and different domains. We have merged classes from hate base to arrive at those domains. *Sexism* contains keywords that belong to *gender* and *sexual orientation*. *Racism* includes classes of *ethnicity* and *nationality*. *Religious Hate* includes words belonging to Religion category on Hatebase⁴. Extremism keywords were handcrafted as described in section 3.2.

English was the dominant language in terms of the total keywords (500) extracted. Besides Hindi, *racism* is the major domain for which all languages have highest number of hate keywords.

3.4. HSMerge Dataset

3.4.1. Dataset Preparation

In order to build the dataset for different tasks with gold labels, we utilize 10 publicly available datasets for different types (tasks) of hate speech detection. We sampled annotated examples from diverse English datasets: GAO hatefulness [5], TRAC aggression [6], offensive [37], racism and sexism [8], MLMA [1], attack [7], CONAN [38], MMHS150k [39].

We map labels from these datasets into different domains: abusive, non-abusive, sexism, racism and religious hate. For OLID [37] *OFFENSIVE* maps to *Abuse*. GAO and WUL had binary

²<https://www.counterextremism.com/daily-dose-archive>

³<https://github.com/twintproject/twint/>

⁴<https://hatebase.org/>

Table 1

Keywords language & domain distribution

Language	Sexism	Racism	Abuse	Religious Hate	Extremism
English	69	418	30	33	88
Hindi	75	16	9	8	-
Arabic	27	35	17	17	-
French	13	85	13	9	-
German	21	67	13	11	-
Spanish	43	77	21	2	-
Total	248	698	103	80	88

Table 2

Tweets distribution

Languages	Raw	Processed
English	184339	105120
Hindi	59321	32734
Arabic	27374	5394
French	75126	20809
German	30868	8631
Spanish	120632	55148
Total	497660	227836

labels, while the original TRAC uses three labels: non-aggressive, covertly-aggressive, and openly-aggressive. We relabel the first as non-abusive, and the other two as abusive. For MLMA we create labels on the basis of target groups: target group *RELIGION* becomes *RELIGIOUS HATE*. For MLMA [1] target group *GENDER* and *SEXUAL ORIENTATION* maps to *SEXISM*, target group *RELIGION* becomes *RELIGIOUS HATE*, target group *ORIGIN* becomes *RACISM*. For HASOC [9] *OFFENSIVE* maps to *ABUSE*. For Conan[40] target groups *MUSLIMS*, *JEWES* becomes *RELIGIOUS HATE* labels, target groups *MIGRANTS*, *POC* maps to *RACISM*, target group *WOMEN*, *LGBT+* maps to *SEXISM* and *DISABLED* maps to general *ABUSE* group. For MMHS150K [39] we got *RACISM*, *SEXISM*, *RELIGIOUS HATE* & *OTHERHATE*.

For some of these data we had to hydrate tweets from Twitter using official Twitter API as only the tweet ids had been provided. For Waseem and Founta [38], majority of the tweets were not available on Twitter which led to less number of tweets compared to the official dataset. We standardized these datasets into a single *HSMerge* dataset. Details of classes and total samples is available in Table 3.

Table 3
HSMerge open source gold dataset

Datasets	Domain	Hate	Non-Hate	Total
OLID	Offensive / NotHate	4400	8840	13240
COnan	Racism/ Sexism/ Religious/ Abuse/ OtherHate	938/1025/1921/175/179	-	4238
TRAC	Hateful / NotHate	5993	4348	10341
MLMA	Racism/ Sexism/ Religious/ OtherHate	2448 / 1152 / 68/1979	-	5647
MLA150K	Racism/ Sexism/ Religious/ OtherHate/ NotHate	11925 / 7365 / 163 / 5811	112845	138209
HASOC	Hateful / NotHate	2261	3591	5852
WASEEM	Racist/Sexist/NotHate	2753/15	7640	10408
WUL	Hate/NotHate	8817	62937	71754
GAO	Hate / NotHate	244	675	909

4. Experiments

4.1. Overall Methodology and Models

In this section, we describe the various components of our pipeline as shown in figure 1 and the models used for semi-supervised annotation process in different settings. Most of the hate speech detection tasks depend on manual annotation that limits the number of samples that can be labeled. We used a hierarchical approach to validate and refine our initial keyword based dataset from Twitter.

1. Monolingual hate speech detection models from Hate-ALERT used as zero shot pipeline to get high level labels of *Hate* or *NoHate* in monolingual settings.
2. Multilingual *mbert* fine-tuned model to get high level labels.
3. Translate the raw data of five languages (other than English) to English.
4. Use multilingual binary classifier (developed in step 2) on translated data to classify tweets at high level.
5. Utilize Google Perspective API to predict toxicity of tweets for high level labels.
6. Use two way voting to predict final binary labels.
7. Fine-tuned "distilbert-base-uncased" model on only the English dataset, and used it to do annotations for the English translations of the data collected from Twitter in 5 languages.
8. To validate our domain specific labels we trained various models on HSMerge data shown in Table 3 and predicted on hateful samples obtained from step 6.

4.2. Experimental Setting

To evaluate the performance of the models, we used weighted average F1 for benchmarking on validation set. All the experiments were done on NVIDIA A100 GPU with up to 20 GiB RAM.

4.3. Hatefulness (Hate/No-Hate)

We utilised hate speech models to perform initial validation on the raw LAHM dataset; details can be found in Section 5.

4.3.1. Monolingual Experiments

We utilized BERT based language specific hate speech models from HuggingFace. The models used were trained, validated and tested on the same language.

We utilized the following language models for our experiments:

1. Hate-speech-CNERG/dehatebert-mono-english⁵
2. l3cube-pune/hate-multi-roberta-hasoc-hindi⁶
3. Hate-speech-CNERG/dehatebert-mono-french⁷
4. Hate-speech-CNERG/dehatebert-mono-german⁸
5. Hate-speech-CNERG/dehatebert-mono-spanish⁹

Binary hate labels distribution results from monolingual experiments are presented in table 8

4.3.2. Machine Translation Experiments

We adopted the methodology to utilize open source machine translation models for translation of English data to multilingual data and vice-versa. To select the best machine translation model for each of our languages, we evaluated a number of models on a small manually annotated dataset. The translations were carried out for each of these languages: Hindi, Arabic, French, German and Spanish. Two types of translations were carried out using the translation models:

1. Translation of English dataset collected from various open source hate speech datasets. All the English language samples in this dataset were translated to other 5 languages.
2. Translation of multilingual data collected from Twitter. This data was collected for the above 5 languages, and each of this language data was translated into corresponding English language data.

Consideration for choosing and evaluating the translation models was based on whether they were open source, free/easy to use, and the translation quality. Translation models used were:

1. Google sheets translation
2. m2m-100-1.2B¹⁰
3. IndicTrans¹¹

The evaluation metrics used were bleu, rougeL and semantic similarity. For semantic similarity, *all-mpnet-base-v2* model from *sentence-transformers* library was used to calculate the cosine score between input and translated sentence embeddings. Figure 4.3.2 shows comparisons of bleu, rougeL and semantic similarity scores for the 3 models on different languages. After comparing the performances and taking other considerations into account, *indicTrans* model was selected for all translations of Indic languages (Hindi), and *m2m-100-1.2B* was selected for the other languages.

⁵<https://huggingface.co/Hate-speech-CNERG/dehatebert-mono-english>

⁶[l3cube-pune/hate-roberta-hasoc-hindi](https://github.com/l3cube-pune/hate-roberta-hasoc-hindi)

⁷<https://huggingface.co/Hate-speech-CNERG/dehatebert-mono-french>

⁸<https://huggingface.co/Hate-speech-CNERG/dehatebert-mono-german>

⁹<https://huggingface.co/Hate-speech-CNERG/dehatebert-mono-spanish>

¹⁰<https://github.com/UKPLab/EasyNMT>

¹¹<https://github.com/AI4Bharat/indicTrans>




Figure 2: bleu scores for model translations




Figure 3: rougeL scores for model translations




Figure 4: Semantic similarity scores for model translations

4.3.3. Binary Multilingual Experiments

We utilized *mbert* which is a pre-trained model on the top 104 languages with the largest Wikipedia corpus using a masked language modeling (MLM) objective, to fine-tune for binary

classification task to perform few shot learning experiments. The aim was to make the classifier learn over samples from a few languages and test over other languages. The training dataset was predominant in English, curated from multiple open source resources. In addition to that we introduced language samples for a few other languages using a neural machine translation based open source model (m2m_100_1.2B)¹². NMT is an approach to achieve machine translation using artificial neural nets to predict the best possible sequence of words.

The model was trained, tested and validated on the custom multilingual dataset curated with the class distribution shown in Table 4 per chosen language. We utilized this model to make predictions on the curated dataset.

Table 4

Multilingual binary classifier training data

Languages	Hate	NoHate	Total
English	20620	18171	38791
Arabic	20539	18138	38680
Hindi	20541	18138	38679
French	20542	18138	38680
Total	82242	72584	154826

4.3.4. Perspective API Experiments

Perspective API is a free to use API that uses machine learning to identify toxicity in the comments. It has API rate limitations of 1 request per second. Due to this limitation we utilized it on all languages except English. We utilized Google Cloud to generate the Perspective API key for sending requests, and specifically *CommentAnalyzer* API to get the toxicity score for each data point. The *CommentAnalyzer* API supports all the 5 languages we have considered. We used the Twitter cleaned data we had collected for the 5 languages to get the Toxicity score to classify each tweet as hate or not-hate. Numbers are reported in Table 8.

4.4. Multi-Domain Experiments

This section describes the multi-domain experiments to validate labels on the LAHM dataset.

4.4.1. Machine Translation Based Multi-Domain Experiments

HSMerge data present in Table 3 was used for the multi-domain classification of hate speech. This dataset was pre-processed by removing URLs, mentions/usernames, null values and duplicates.

The aim was to utilize transfer learning to fine-tune a model on only the English dataset, and use it to do annotations for the English translations of the data collected from Twitter in other 5 languages.

¹²<https://github.com/pytorch/fairseq/tree/main/examples/m2m100>

The model used for fine-tuning on this dataset was "*distilbert-base-uncased*"¹³. DistilBERT is a distilled version of the BERT base model. It has 40 percent less parameters than bert-base-uncased, is 60 percent faster and retains almost 97 percent of BERT’s performances¹⁴. A linear layer on top of the pooled output of the model was added for the multi-class classification. During training, the maximum sequence length was limited to 256, batch size was set to 32, learning rate 5e-5 and was trained for 4 epochs. The dataset was split into train, validation and test set using stratified random sampling. Dataset details :

1. 20,174 training samples
2. 5,044 validation samples
3. 6,305 test samples

The model evaluation metrics on the test set are given in Table 5.

Table 5

Translated multiclass metrics

Domain	precision	recall	f1-score	support
racism	0.94	0.95	0.94	2882
sexism	0.93	0.93	0.93	2295
religious hate	0.92	0.93	0.92	409
abuse	0.94	0.91	0.93	589
extremism	0.83	0.82	0.82	130
accuracy			0.93	6305
macro avg	0.91	0.91	0.91	6305
weighted avg	0.93	0.93	0.93	6305

4.4.2. Cross lingual based Multi-Domain Experiments

To validate our fine-grained labels with keywords from hatebase, we experimented with cross-lingual models trained on HSmerge dataset as shown in Table 3. In this experiment setting we fine tuned the *mbert* model for the multi-class classification task and trained the model on HSmerge data in Table 3 and do the prediction on the other languages. Empirical results can be found under Table 7. In this zero-shot setting, no other language samples were given to the model after fine-tuning on the English dataset and also machine translation was not involved at all in either the pre-training or fine-tuning.

We couldn’t manage to maintain the class distribution, and the *extremism* class was left with comparatively less data points. To counter the issue we introduced custom class weights to provide equal attention to the minority class which in our case was *extremism*. This was done using a weighted class random sampler *WeightedSampler* for class imbalance, so that all classes have equal probability. The empirical scores achieved by the cross-lingual model can be found in the Table 7.

¹³<https://huggingface.co/distilbert-base-uncased>

¹⁴<https://arxiv.org/abs/1910.01108>

The open source *mbert* model we fine-tuned for our use case has previously been tested for zero-shot experimentation¹⁵ and did manage to achieve decent empirical scores shown in Table 6. Dataset details :

1. 19298 training samples
2. 6432 validation samples
3. 6434 test samples

Table 6

Bert zero shot gold **F1** scores

Model	English	Spanish	German	Arabic
BERT-Zero Shot	81.4	74.3	70.5	62.1

Table 7

HSMerge cross-lingual metrics

Domains	precision	recall	f1-score	support
racism	0.93	0.95	0.94	2161
sexism	0.93	0.91	0.92	1722
religious hate	0.93	0.93	0.93	307
abuse	0.93	0.92	0.92	442
extremism	0.85	0.80	0.83	97
accuracy			0.93	4729
macro avg	0.91	0.90	0.91	4729
weighted avg	0.93	0.93	0.93	4729

5. Evaluation

We performed initial level of validation (Hatefulness) experiments on the raw LAHM dataset using open source monolingual models, Perspective API and the multilingual binary classifier we trained. We compared the results for monolingual and multilingual and noticed that our multilingual few-shot learning based classifier out-performed the open source monolingual BERT based models from HuggingFace except on Arabic and French languages where the latter did a better job; details can be found in Table 8. Based on the analysis we took at least 2 votes from the models which performed better, details shown in Table 9.

For Hindi, Perspective API performed better by predicting 31 percent as hate compared to 12 percent for binary classifier. For Arabic, the binary classifier performed significantly better.

For domain validation we utilised the zero-shot cross-lingual, multi-domain and machine translation based multi-domain models trained with HSmerge data in Table 3 and performed

¹⁵<https://github.com/google-research/bert/blob/master/multilingual.md#results>

Table 8

All models prediction on high level label Hate and NoHate

Language	Monolingual		Multilingual		Perspective	
	Hate	NoHate	Hate	NoHate	Hate	NoHate
English	7247	101345	11380	99598	-	-
Hindi	210	260	3970	7247	1674	3720
Arabic	6469	26265	3914	28820	1376	31358
French	9148	24185	8893	24440	7540	13269
German	468	23910	3905	20473	4467	4164
Spanish	2088	10087	4734	7441	24952	30196
Total	25630	186052	36796	188019	40009	82707

Table 9

Voting across multiple models for hate detection

Language	Mono-Lingual	Multi-lingual	Perspective API
English	✓		✓
Hindi		✓	✓
Arabic	✓	✓	
French	✓	✓	
German		✓	✓
Spanish	✓		✓

Table 10

Cross-lingual multi-domain predictions distribution

Language	Abuse	Sexism	Racism	Religious	Extremism
English	4228	2934	1566	751	154
Arabic	1895	994	736	164	5631
Hindi	1511	440	282	22	1715
French	813	779	2393	750	4413
German	180	123	30	4	131
Spanish	818	483	139	11	637
Total	9445	5753	5146	1702	12681

Table 11
Machine Translation multi-domain predictions distribution

Language	Abuse	Sexism	Racism	Religious	Extremism
English	8691	6693	2495	261	524
Arabic	430	1302	209	259	1138
Hindi	132	204	67	137	105
French	745	1004	545	273	538
German	232	608	186	58	98
Spanish	2082	2024	1519	347	1310
Total	12312	11835	5021	1335	3713

predictions on the different languages. Label distribution from cross-lingual models are shown in Table 10. Label distribution of predictions for different languages on the LAHM dataset from multi-class translated model is in Table 11.

For Hindi, 31.6 percent samples were Abuse, while the lowest were 10.3 percent for Racism. For Arabic, Sexism and Extremism contributed to 73 percent of all hate samples. For French, 32 percent were Abuse and lowest 8.3 percent were Religious Hate. For German, 51 percent were Sexism while the lowest were Religious Hate and Extremism together amounting to 13 percent. For Spanish, Abuse and Sexism had 27 percent each, while Religious Hate had 5 percent of samples.

6. Conclusion

We have presented the LAHM dataset, a large scale semi-supervised training dataset for multilingual and multi-domain hate speech identification, we created by using 3 layer annotation pipeline and combination of monolingual, multilingual and cross-lingual models. To the best of our knowledge, LAHM is the largest of its kind, containing close to 300k tweets across 6 languages and 5 domains.

LAHM enables cross-lingual abusive language detection across five domains and in-depth interplay between language shift and domain shift. We have profiled LAHM as a comprehensive resource for evaluating hate speech detection through a series of cross-domain experiments in monolingual, multilingual and cross-lingual setups with state of the art transfer learning models.

We hope that LAHM will inspire more efforts in understanding and building semi-supervised large scale multilingual and multi-domain abusive language detection datasets.

7. Future Work

For future work we explored leetspeak detection and identification on social networks. A lot of hate content on social media uses leetspeak to evade moderators and automated systems. We collected hate keywords belonging in this category for each of the 5 domains, and experimented with a set of leets for each keyword to extract the leetspeak hate content. We plan to use this for future work in multilingual and multi-domain settings.

References

- [1] N. Ousidhoum, Z. Lin, H. Zhang, Y. Song, D.-Y. Yeung, Multilingual and multi-aspect hate speech analysis, in: Proceedings of EMNLP, Association for Computational Linguistics, 2019.
- [2] Y.-L. Chung, E. Kuzmenko, S. S. Tekiroglu, M. Guerini, Conan-counter narratives through nichesourcing: a multilingual dataset of responses to fight online hate speech, arXiv preprint arXiv:1910.03270 (2019).
- [3] T. Davidson, D. Warmesley, M. Macy, I. Weber, Automated hate speech detection and the problem of offensive language, in: Proceedings of the International AAAI Conference on Web and Social Media, volume 11, 2017.
- [4] V. Kolhatkar, M. Taboada, A corpus for the analysis of online news comments (????).
- [5] L. Gao, R. Huang, Detecting online hate speech using context aware models, arXiv preprint arXiv:1710.07395 (2017).
- [6] R. Kumar, A. K. Ojha, S. Malmasi, M. Zampieri, Benchmarking aggression identification in social media, in: Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018), 2018, pp. 1–11.
- [7] E. Wulczyn, N. Thain, L. Dixon, Ex machina: Personal attacks seen at scale, in: Proceedings of the 26th international conference on world wide web, 2017, pp. 1391–1399.
- [8] Z. Waseem, D. Hovy, Hateful symbols or hateful people? predictive features for hate speech detection on twitter, in: Proceedings of the NAACL Student Research Workshop, Association for Computational Linguistics, San Diego, California, 2016, pp. 88–93. URL: <http://www.aclweb.org/anthology/N16-2013>.
- [9] T. Mandl, S. Modha, C. Mandlia, D. Patel, A. Patel, M. Dave, Hasoc-hate speech and offensive content identification in indo-european languages, 2019.
- [10] A. Basile, C. Rubagotti, Crotonemilano for ami at evalita2018. a performant, cross-lingual misogyny detection system., EVALITA Evaluation of NLP and Speech Tools for Italian 12 (2018) 206.
- [11] M. Wiegand, M. Siegel, J. Ruppenhofer, Overview of the germeval 2018 shared task on the identification of offensive language (2018).
- [12] J. Salminen, M. Hopf, S. A. Chowdhury, S.-g. Jung, H. Almerexhi, B. J. Jansen, Developing an online hate classifier for multiple social media platforms, Human-centric Computing and Information Sciences 10 (2020) 1–34.
- [13] S. A. Chowdhury, H. Mubarak, A. Abdelali, S.-g. Jung, B. J. Jansen, J. Salminen, A multi-

- platform arabic news comment dataset for offensive language detection, in: Proceedings of the 12th Language Resources and Evaluation Conference, 2020, pp. 6203–6212.
- [14] M. Wiegand, J. Ruppenhofer, A. Schmidt, C. Greenberg, Inducing a lexicon of abusive words—a feature-based approach (2018).
- [15] M. Karan, J. Šnajder, Cross-domain detection of abusive language online, in: Proceedings of the 2nd workshop on abusive language online (ALW2), 2018, pp. 132–137.
- [16] E. W. Pamungkas, V. Patti, Cross-domain and cross-lingual abusive language detection: A hybrid approach with deep learning and a multilingual lexicon, in: Proceedings of the 57th annual meeting of the association for computational linguistics: Student research workshop, 2019, pp. 363–370.
- [17] E. W. Pamungkas, V. Basile, V. Patti, Misogyny detection in twitter: a multilingual and cross-domain study, *Information Processing & Management* 57 (2020) 102360.
- [18] J. S. Meyer, B. Gambäck, A platform agnostic dual-strand hate speech detector, in: ACL 2019 The Third Workshop on Abusive Language Online Proceedings of the Workshop, Association for Computational Linguistics, 2019.
- [19] K. Wang, D. Lu, S. C. Han, S. Long, J. Poon, Detect all abuse! toward universal abusive language detection models, *arXiv preprint arXiv:2010.03776* (2020).
- [20] A. Arango, J. Pérez, B. Poblete, Hate speech detection is not as easy as you may think: A closer look at model validation, in: Proceedings of the 42nd international acm sigir conference on research and development in information retrieval, 2019, pp. 45–54.
- [21] Z. Waseem, J. Thorne, J. Bingel, Bridging the gaps: Multi task learning for domain transfer of hate speech detection, in: *Online harassment*, Springer, 2018, pp. 29–55.
- [22] M.-A. Rizoii, T. Wang, G. Ferraro, H. Suominen, Transfer learning for hate speech detection in social media, *arXiv preprint arXiv:1906.03829* (2019).
- [23] T. Caselli, V. Basile, J. Mitrović, M. Granitzer, Hatebert: Retraining bert for abusive language detection in english, *arXiv preprint arXiv:2010.12472* (2020).
- [24] A. Koufakou, E. W. Pamungkas, V. Basile, V. Patti, Hurtbert: incorporating lexical features with bert for the detection of abusive language, in: Proceedings of the fourth workshop on online abuse and harms, 2020, pp. 34–43.
- [25] G. Glavaš, M. Karan, I. Vulic, Xhate-999: Analyzing and detecting abusive language across domains and languages, *Association for Computational Linguistics*, 2020.
- [26] M. Mozafari, R. Farahbakhsh, N. Crespi, Hate speech detection and racial bias mitigation in social media based on bert model, *PloS one* 15 (2020) e0237861.
- [27] K. B. Ozler, K. Kenski, S. Rains, Y. Shmargad, K. Coe, S. Bethard, Fine-tuning bert for multi-domain and multi-label incivil language detection, in: Proceedings of the Fourth Workshop on Online Abuse and Harms, 2020, pp. 28–33.
- [28] S. S. Aluru, B. Mathew, P. Saha, A. Mukherjee, A deep dive into multilingual hate speech classification, in: *Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track: European Conference, ECML PKDD 2020, Ghent, Belgium, September 14–18, 2020, Proceedings, Part V*, Springer International Publishing, 2021, pp. 423–439.
- [29] N. Vashistha, A. Zubiaga, Online multilingual hate speech detection: experimenting with hindi and english social media, *Information* 12 (2021) 5.
- [30] M. O. Ibrohim, I. Budi, Translated vs non-translated method for multilingual hate speech

- identification in twitter, *Int. J. Adv. Sci. Eng. Inf. Technol* 9 (2019) 1116–1123.
- [31] M. Corazza, S. Menini, E. Cabrio, S. Tonelli, S. Villata, Hybrid emoji-based masked language models for zero-shot abusive language detection, in: *Findings of the Association for Computational Linguistics: EMNLP 2020*, Association for Computational Linguistics, Online, 2020, pp. 943–949. URL: <https://aclanthology.org/2020.findings-emnlp.84>. doi:10.18653/v1/2020.findings-emnlp.84.
- [32] H. Ahn, J. Sun, C. Y. Park, J. Seo, NLPDove at SemEval-2020 task 12: Improving offensive language detection with cross-lingual transfer, in: *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, International Committee for Computational Linguistics, Barcelona (online), 2020, pp. 1576–1586. URL: <https://aclanthology.org/2020.semeval-1.206>. doi:10.18653/v1/2020.semeval-1.206.
- [33] J. M. Pérez, A. Arango, F. Luque, ANDES at SemEval-2020 task 12: A jointly-trained BERT multilingual model for offensive language detection, in: *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, International Committee for Computational Linguistics, Barcelona (online), 2020, pp. 1524–1531. URL: <https://aclanthology.org/2020.semeval-1.199>. doi:10.18653/v1/2020.semeval-1.199.
- [34] L. Stappen, F. Brunn, B. Schuller, Cross-lingual zero-and few-shot hate speech detection utilising frozen transformer language models and axel, *arXiv preprint arXiv:2004.13850* (2020).
- [35] T. Dadu, K. Pant, Team rouges at SemEval-2020 task 12: Cross-lingual inductive transfer to detect offensive language, in: *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, International Committee for Computational Linguistics, Barcelona (online), 2020, pp. 2183–2189. URL: <https://aclanthology.org/2020.semeval-1.290>. doi:10.18653/v1/2020.semeval-1.290.
- [36] T. Ranasinghe, M. Zampieri, Multilingual offensive language identification with cross-lingual embeddings, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Online, 2020, pp. 5838–5844. URL: <https://aclanthology.org/2020.emnlp-main.470>. doi:10.18653/v1/2020.emnlp-main.470.
- [37] M. Zampieri, P. Nakov, S. Rosenthal, P. Atanasova, G. Karadzhov, H. Mubarak, L. Derczynski, Z. Pitenis, Ç. Çöltekin, Semeval-2020 task 12: Multilingual offensive language identification in social media (offenseval 2020), *arXiv preprint arXiv:2006.07235* (2020).
- [38] A. M. Founta, C. Djouvas, D. Chatzakou, I. Leontiadis, J. Blackburn, G. Stringhini, A. Vakali, M. Sirivianos, N. Kourtellis, Large scale crowdsourcing and characterization of twitter abusive behavior, in: *Twelfth International AAAI Conference on Web and Social Media*, 2018.
- [39] R. Gomez, J. Gibert, L. Gomez, D. Karatzas, Exploring hate speech detection in multimodal publications, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2020, pp. 1470–1478.
- [40] Fanton, Margherita and Bonaldi, Helena and Tekiroğlu, Serra Sinem and Guerini, Marco, Human-in-the-Loop for Data Collection: a Multi-Target Counter Narrative Dataset to Fight Online Hate Speech, in: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, 2021.