

# K-MHaS: A Multi-label Hate Speech Detection Dataset in Korean Online News Comment

Jean Lee<sup>1</sup> Taejun Lim<sup>1</sup> Heejun Lee<sup>3</sup> Bogeun Jo<sup>3</sup>  
Yangsook Kim<sup>4</sup> Heegeun Yoon<sup>5</sup> Soyeon Caren Han<sup>1,2\*</sup>

The University of Sydney<sup>1</sup> The University of Western Australia<sup>2</sup>  
BigWave AI<sup>3</sup> Keimyung University<sup>4</sup> National Information Society Agency<sup>5</sup>  
{jean.lee, tlim6535, caren.han}@sydney.edu.au  
{leej3471, iambgjo}@gmail.com  
yangsok.kim@kmu.ac.kr hgyoon@nia.or.kr

## Abstract

Online hate speech detection has become an important issue due to the growth of online content, but resources in languages other than English are extremely limited. We introduce K-MHaS<sup>1</sup>, a new multi-label dataset for hate speech detection that effectively handles Korean language patterns. The dataset consists of 109k utterances from news comments and provides a multi-label classification using 1 to 4 labels, and handles subjectivity and intersectionality. We evaluate strong baselines on K-MHaS. KR-BERT with a sub-character tokenizer outperforms others, recognizing decomposed characters in each hate speech class.

## 1 Introduction

The growth of online content including social media (Zampieri et al., 2020), news comments (Gao and Huang, 2017), Wikipedia (Wulczyn et al., 2017), and in-game chat (Weld et al., 2021) presents challenges in detecting hate speech using advanced Natural Language Processing. Hate speech is language that attacks or diminishes individuals or groups based on certain characteristics such as physical appearance, religion, gender, or other attributes, and it can occur across different linguistic styles (Fortuna and Nunes, 2018). Hate speech detection is intrinsically a complex task (Wang et al., 2020) due to the fuzzy boundary with other overlapping concepts such as abusive language (Nobata et al., 2016), toxic comments (Wulczyn et al., 2017), or offensive language (Davidson et al., 2017).

Recently, the rise in popularity of Korean TV, movies, and music (e.g. Squid Game, BTS) has led to many young people showing an interest in learning Korean. This phenomenon could result in exposure to harmful content and hate speech in

Korean. However, (1) the most common language in hate speech research is English and only limited resources are available in other languages such as Arabic (Mubarak et al., 2017), Dutch (Caselli et al., 2021), and Korean (Moon et al., 2020). In addition, most datasets are annotated (2) using a single label classification of particular aspects, even though the subjectivity of hate speech cannot be explained with a mutually exclusive annotation scheme.

We propose K-MHaS, a Korean multi-label hate speech detection dataset that allows overlapping labels associated with intersectionality, a concept from sociology that identifies combined attributes (Crenshaw, 1989). Our dataset consists of 109,692 utterances from Korean online news comments, labeled with 8 fine-grained hate speech classes. K-MHaS is compatible with previous work on hate speech in other languages, by providing binary classification and multi-label classification from 1(one) to 4(four) labels.

We investigate the K-MHaS dataset by analyzing label distribution, keywords, and label pairs. In addition, we provide strong baseline pre-trained language models using Multilingual-BERT, KoELECTRA, KoBERT, and KR-BERT, and compare the results using six metrics for multi-label classification tasks. Overall, the KoELECTRA model achieves the best performance for all labels, indicating the effects of the pre-training data source. The KR-BERT with a sub-character-level tokenizer outperforms the others on several label pairs, showing that decomposing various Korean characters is essential for the task. Our contribution can be summarized as follows:

- We propose a large size Korean multi-label hate speech detection dataset that represents Korean language patterns effectively;
- We propose a multi-label hate speech annotation scheme, which can handle the subjectivity of hate speech and the intersectionality;

\*Corresponding author (caren.han@sydney.edu.au)

<sup>1</sup>The dataset is available at <https://github.com/adlnlp/K-MHaS>.

Publication	Language	Source	Data size	Labels	M-label
Waseem and Hovy (2016)	English	Twitter	16.2k	Sexism, Racism, Neither	N
Davidson et al. (2017)	English	Twitter	24.8k	Hate Speech, Offensive, Neither	N
Wulczyn et al. (2017)	English	Wikipedia comments	115k	Toxic, Severe Toxic, Obscene, Threat, Insult, Identity Hate, Neutral	Y
Ibrohim and Budi (2019)	Indonesian	Twitter	11k	(a) Individual, Group (b) Religion, Race, Physical, Gender, Other (c) Weak, Moderate, Strong Hate Speech	P
Fortuna et al. (2019)	Portuguese	Twitter	5.6k	(a) Hate Speech, Not Hate Speech (b) Sexism, Body, Origin, Homophobia, Racism, Ideology, Religion, Health, Other-Lifestyle	P
Ousidhoum et al. (2019)	English French Arabic	Twitter	6k (EN) 4k (FR) 3k (AR)	Labels for five different aspects (a) Directness, (b) Hostility, (c) Target, (d) Group, and (e) Annotator	P
Moon et al. (2020)	Korean	News comments	9k	(a) Hate Speech, Offensive, None (b) Gender, Others, None	N
<b>Ours</b>	<b>Korean</b>	<b>News comments</b>	<b>109k</b>	<b>(a) Hate Speech, Not Hate Speech (b) Politics, Origin, Physical, Age, Gender Religion, Race, Profanity, Not Hate Speech</b>	<b>Y</b>

Table 1: Comparison of datasets. A "M-label" indicates a multi-label annotation scheme that allows overlapping labels for intersectionality (P = partially applied). The (a) - (e) indicates a layer containing a single label from each aspect.

- We evaluate strong baseline experiments on our dataset using Korean-BERT-based language models with six different metrics.

## 2 Korean Multi-label Hate Speech Detection Dataset (K-MHaS)

Our dataset is based on the Korean online news comments available on Kaggle<sup>2</sup> and Github<sup>3</sup>. The unlabeled raw data was collected between January 2018 and June 2020. In order to curate the data, we randomly select more than 109,692 news comments. Our data preprocessing is designed to tokenize a Korean character and to filter the length. We remove URLs and bad characters (e.g. U+1100 to U+11FF - Hangeul Jamo) using regular expressions while keeping uppercase and lowercase letters in English and emoji. We discard sentences with fewer than 10 characters as it is often only one word. For the data derived from online comments, we normalized repeated characters by truncating their number of consecutive repetitions to two.

**Multi-label Annotation** We consider a multi-label annotation scheme in order to deliver fine-grained hate speech categories and intersectionality from the overlapping labels. The annotation scheme has two layers: (a) binary classification ('Hate Speech' or 'Not Hate Speech') and (b) fine-grained classification (8 labels or 'Not Hate

Speech'). For the fine-grained classification, a 'Hate Speech' class from the binary classification is broken down into 8 classes associated with the hate speech category<sup>4</sup>. As shown in Table 1, this scheme allows non-exclusive concepts, accounting for the overlapping shades of given categories. We select the 8 hate speech classes in order to reflect the social and historical context as the nature of hate speech is different in each language (Kang et al., 2020). For example, the 'politics' class is chosen due to a significant influence on the style of Korean hate speech.

**Annotation Instructions** Given the subjectivity of the task and our annotation scheme, we perform a preliminary round to identify the topics of hate speech and develop annotation instructions. We begin with the common categories of hate speech found in literature and match the keywords for each category. After the preliminary round, we investigate the results to merge or remove labels in order to provide the most representative subtype labels of hate speech contextual to the cultural background. Our annotation instruction includes the criteria as follows: **Politics**: hate speech based on political stance; **Origin**: hate speech based on place of origin or identity; **Physical**: hate speech based on physical appearance (e.g. body, face) or disability; **Age**: hate speech based on age; **Gender**: hate

<sup>2</sup><https://www.kaggle.com/datasets/junbumlee/kcbert-pretraining-corpus-korean-news-comments>

<sup>3</sup><https://github.com/kocohub/korean-hate-speech>

<sup>4</sup> Fine-grained labels (matching in Korean) : Politics (정치성향차별), Origin (출신차별), Physical (외모차별), Age (연령차별), Gender (성차별), Religion (종교차별), Race (인종차별), and Profanity (혐오욕설)




Figure 1: Overview of Annotation Process.

speech based on gender or sexual orientation (e.g. woman, homosexual); **Religion**: hate speech based on religion; **Race**: hate speech based on ethnicity; **Profanity**: hate speech in the form of swearing, cursing, cussing, obscene words, or expletives; or an unspecified hate speech category from above; and **Not Hate Speech**.

Our annotation instructions explain a two-layered annotation to (a) distinguish hate and not hate speech, and (b) the categories of hate speech. Annotators are requested to consider given keywords or alternatives of each category within social, cultural, and historical circumstances. For example, a comment using the word “women” is not hate speech, whereas, if it is critical of “women” or uses language that attacks the group, it is classified as ‘gender’. Notably, we annotate multi-labels if a comment includes several hate speech categories. Since hate speech can be varied, any comments in the form of swearing or cursing are marked as ‘profanity’. For instance, a comment containing hate speech about appearance, political stance, and gender in profane language (e.g. “fuck you ugly communist bitch.”)<sup>5</sup> is labeled within ‘physical’, ‘politics’, ‘gender’ and ‘profanity’ classes.

**Annotation Process** Five native speakers were recruited for manual annotation in both the preliminary and main rounds. During the preliminary round, we facilitated the annotation instructions by conducting an annotators’ discussion and providing some examples of keywords for each class. As shown in Figure 1, we introduced an iterative process that enables faster annotation in the main round. We provided an ‘uncertain’ additional field that was used for the unspecified label in annotation guidelines or when the annotator had difficulties in choosing labels. Any ‘uncertain’ labeled data was flagged by individual annotators, then reviewed by five annotators. The final labels were chosen based on the majority vote, and the annotation guidelines were updated to handle similar cases. Additionally, the other labeled data was reviewed, in line

<sup>5</sup> (Korean) “면상도 개 조가치 생겼네 개빨갱이년”

Label Types	Count (%)	
Total Utterances	<b>109,692 (100%)</b>	
Multi-label (Hate Speech)	1 label (Single)	36,470 (33.2%)
	2 labels	12,073 (11.0%)
	3 labels	1,440 (1.3%)
	4 labels	94 (0.1%)
Not Hate Speech	59,615 (54.3%)	

Table 2: Dataset Statistics. The total is the combination of all ‘hate speech’ and ‘not hate speech’ label. Together the ‘hate speech’ label makes up 45.7% of the data.

Class	Count - Single (%)	Count - Multi (%)
Politics	6,931 (19.0%)	4,961 (17.2%)
Origin	5,739 (15.7%)	4,458 (15.5%)
Physical	5,443 (14.9%)	3,364 (11.7%)
Age	4,192 (11.5%)	3,178 (11.0%)
Gender	3,348 (9.2%)	4,696 (16.3%)
Religion	1,862 (5.1%)	513 (1.8%)
Race	160 (0.4%)	163 (0.6%)
Profanity	8,795 (24.1%)	7,509 (26.0%)

Table 3: Fine-grained label distributions on *hate speech* labels. A ‘not hate speech’ label is not included. A single means 1 label and a multi is the sum of 2, 3, and 4 labels. A multi-labeled data counts each overlapping class.

with the annotation guideline by two random annotators for the final dataset. The inter-annotator agreement returns an average Cohen Kappa score of 0.892, indicating substantial agreement (Siegel and Castellan, 1988).

### 3 Dataset Analysis

K-MHaS dataset contains 109,692 comments as shown in Table 2. For binary classification, the proportion of the ‘hate speech’ (45.7%) and ‘not hate speech’ (54.3%) satisfies data balancing. The ‘hate speech’ label consists of a single label (33.2%) and multi-labels (12.4%), containing from 2 to 4 labels. Other hate speech datasets reviewed have an approximate ratio of ‘hate speech’ to ‘not hate speech’ of around 40% (Vidgen and Derczynski, 2020). Our dataset is consistent with this figure, where the ‘hate speech’ in a single label to ‘not hate speech’ ratio is 38%.

**Label Distribution** Table 3 shows the fine-grained label distribution across our K-MHaS. For both single (s) and multi-label (m) distribution, the ‘profanity’ class (24.1%-s, 26.0%-m) is more frequent than any other class, indicating that swear words are critical for detecting hate speech. Also, the ‘religion’ (5.1%-s, 1.8%-m) and ‘race’ (0.4%-s, 0.6%-m) classes are the smallest portions in both distributions, which are significantly more com-

Rank	Politics	Origin	Physical	Age
1	재앙 (1427)	짱개 (615)	얼굴 (962)	틀 (1918)
2	문재인 (951)	전라도 (596)	돼지 (772)	나이 (599)
3	좌파 (464)	중국 (539)	여자 (294)	노인 (139)
4	좌빨 (402)	쪽 (448)	성형 (216)	충 (112)
5	빨갱이 (367)	짱 (446)	관상 (183)	놈 (106)

Rank	Gender	Religion	Race	Profanity
1	여자 (1704)	개독교 (526)	흑인 (44)	새끼 (1103)
2	남자 (990)	신천지 (460)	백인 (32)	년 (1014)
3	페미 (172)	사이비 (409)	양키 (32)	지랄 (564)
4	맘충 (138)	종교 (305)	감둥이 (19)	개 (459)
5	여성 (134)	예수 (227)	놈 (13)	놈 (404)

Table 4: Top 5 keywords associated with each fine-grained label. The number in brackets is the token count. The keyword analysis is from the total dataset and different from some examples in annotation guidelines.




Figure 2: Average utterance length. (a) label types from 1 to 4 labels. 8 class types (b) in a single label and (c) in multi-labels.

mon in other hate speech datasets. This difference could be because Korea is a highly homogenous monoculture with little variation in race and religion (Kang et al., 2020). Interestingly, the ‘gender’ class (16.3%) occurs at almost twice the frequency in a multi-label distribution, compared to a single label distribution (9.2%). This indicates that gender-based hate speech is used extensively in combined aspects.

**Keyword Analysis** To understand the lexical aspects, we list the top 5 keywords for each hate speech category in Table 4, identifying which tokens are highly associated with each class. In the ‘politics’ class, we find that far-right extremism is dominant, and new tokens such as “catastrophe” [jae ang](재앙) appears related to the former president’s given name ([jae in]) as the two words are near-homophones. Across all classes, one-word tokens are often used in their stem form to modify the meanings of other words. For example, a token [teul] (틀) comes from the word “denture” [teulni] (틀니) which is used as an offensive reference to the elderly. In addition, one-word tokens can be used as a prefix (e.g. “dog” [gae] (개)) or a suffix (e.g. “insect” [chung] (충)), and combined with other neutral words to create a new offensive term.

**Label Pair Analysis** Figure 2 shows the average length of utterance by label count and class type.

Model	F1 (macro)	F1 (micro)	F1 (weighted)	E.M.	AUC	H.L. (↓)
BERT	0.6912	0.8139	0.8119	0.7579	0.8878	0.0464
KoELECTRA	0.7245	<u>0.8493</u>	<b>0.8480</b>	<b>0.7994</b>	<b>0.9122</b>	<u>0.0380</u>
KoBERT	<b>0.7651</b>	0.8413	0.8424	<u>0.7926</u>	<u>0.9083</u>	0.0401
KR-BERT-c	<u>0.7444</u>	<b>0.8500</b>	<u>0.8470</u>	0.7901	0.9028	<b>0.0368</b>
KR-BERT-s	0.7245	0.8445	0.8437	0.7825	0.9076	0.0390

Table 5: Overall multi-label classification performance on K-MHAs for the five baseline models at epoch 4 (E.M.:Exact Match, H.L.:Hamming Loss / KR-BERT-\*: c = character-level, s = sub-character-level)

Metric	# labels	BERT	KoELECTRA	KoBERT	KR-BERT-c	KR-BERT-s
F1 (micro)	1	0.8190	<u>0.8490</u>	0.8320	<b>0.8553</b>	0.8392
	2	0.8043	0.8612	<b>0.8854</b>	0.8405	<u>0.8703</u>
	3	0.7517	0.7987	<u>0.8290</u>	0.7827	<b>0.8329</b>
	4	0.7093	0.7044	0.6832	<u>0.7439</u>	<b>0.7711</b>

Table 6: A breakdown of F1 for multi-label classification from 1 to 4 labels.

The total average length of an utterance is 33 tokens. An increase in the number of labels shows an increasing trend in utterance length, indicating that multi-labeled hate speech contains more linguistic content. The ‘gender’ class has relatively longer lengths (43 tokens) compared to other classes in a single label, whereas all multi-labels utterances have a similar length. This indicates that the gender class has different linguistic features.

## 4 Experiment Setup

**Data Preparation** We split the data into train/test in the proportions of 0.8/0.2. From the training set, we randomly select 0.1 as a validation set (78,977/8,776/21,939 samples for train/val/test sets, preserving the class proportion). The data passed to the models is the preprocessed sentences and binary label vectors.

**Baselines** We select four baselines. 1) **Multi-BERT** (Wolf et al., 2020) is pre-trained on Wikipedia in 104 different languages. We adopted the BERT-Base, uses the WordPiece tokenizer and contains 110M parameters and 119K vocabs. 2) **KoELECTRA** (Park, 2020) is pre-trained on 34GB Korean news, Korean Wikipedia, Namuwiki (Korean-based wiki) and Modu (Korean corpus data publicly provided by the Korean government). The KoELECTRA-Small-v3 is used with the WordPiece tokenizer and contains 14M parameters and 35K vocabs. 3) **KoBERT** (SKTBrain, 2019) is pre-trained on 54M words from Korean Wikipedia, used the SentencePiece tokenizer, 92M parameters and 8K vocabs. 4) **KR-BERT** (Lee et al., 2020) is pre-trained on 2.47GB corpus with 233M words from Korean Wikipedia and news. We applied



Label Pairs	# pairs	F1 (macro)					F1 (micro)				
		BERT	KoELECTRA	KoBERT	KR-BERT-c	KR-BERT-s	BERT	KoELECTRA	KoBERT	KR-BERT-c	KR-BERT-s
<b>Overall Performance (F1)</b>		0.6912	0.7245	<b>0.7651</b>	0.7444	0.7245	0.8139	<u>0.8493</u>	0.8413	<b>0.8500</b>	0.8445
<i>Profanity &amp; Politics</i>	323	0.1959	<u>0.2045</u>	<b>0.2072</b>	0.2013	0.2034	0.8379	<u>0.8853</u>	<b>0.9010</b>	0.8687	0.8616
<i>Profanity &amp; Physical</i>	311	0.1931	0.2061	<u>0.2115</u>	0.2099	<b>0.2121</b>	0.8393	0.9096	0.9331	<b>0.9369</b>	<u>0.9334</u>
<i>Profanity &amp; Origin</i>	269	0.1887	<u>0.1989</u>	0.1987	0.1961	<b>0.2050</b>	0.8144	<u>0.8731</u>	0.8729	0.8661	<b>0.9070</b>
<i>Gender &amp; Origin</i>	242	0.2035	0.2017	<u>0.2134</u>	0.1905	<b>0.2141</b>	0.8920	0.8780	<u>0.9440</u>	0.8354	<b>0.9494</b>
<i>Politics &amp; Origin</i>	224	0.1962	0.1976	<u>0.1991</u>	0.1872	<b>0.2013</b>	0.8666	0.8714	<u>0.8846</u>	0.8295	<b>0.8918</b>
<i>Age &amp; Politics</i>	222	0.1996	<b>0.2114</b>	<u>0.2104</u>	0.1964	0.2014	0.8765	<u>0.9329</u>	<b>0.9357</b>	0.8734	0.8878
<i>Gender &amp; Profanity</i>	181	0.1895	<u>0.1991</u>	0.1957	0.1911	<b>0.2054</b>	0.8157	<u>0.8715</u>	0.8542	0.8450	<b>0.8994</b>
<i>Gender &amp; Physical</i>	177	0.1160	0.1833	<b>0.1958</b>	0.1813	<u>0.1953</u>	0.4562	0.7867	<u>0.8455</u>	0.8045	<b>0.8585</b>
<i>Age &amp; Profanity</i>	132	0.1908	<u>0.2102</u>	<b>0.2139</b>	0.2063	0.2043	0.8414	<u>0.9240</u>	<b>0.9459</b>	0.9105	0.9095
<i>Gender &amp; Age</i>	130	0.1738	<u>0.1781</u>	<b>0.1903</b>	0.1517	0.1339	0.7277	<u>0.7452</u>	<b>0.8159</b>	0.6368	0.5686

Table 7: F1 score for the top 10 two-label pairs on the K-MHaS dataset for the five pre-trained language models at epoch 4 (# total label pairs = 2,439 / KR-BERT-\*: c = character-level tokenizer, s = sub-character-level tokenizer).

either (1) the character-level tokenizer or (2) the sub-character-level tokenizer<sup>6</sup>.

**Evaluation Metrics** In multi-label classification, the prediction contains a set of labels, which means the prediction can be fully correct, partially correct, or fully incorrect. We propose to use the widely used six metrics (Godbole and Sarawagi, 2004) for conducting our multi-label classification, including F1-[macro, micro, weighted], Exact Match, AUC and Hamming Loss (Sorower, 2010).

## 5 Results

**Evaluation for All Labels** The overall performance for all labels is provided in Table 5. The F1(micro) range between 0.8139 (Multi-BERT), 0.8493 (KoELECTRA) and 0.8500 (KR-BERT-c), while the F1(macro) scores show a range from 0.6912 (Multi-BERT) to 0.7651 (KoBERT) with 4 epochs. We observe that all baselines achieve a similar performance, whereas Multi-BERT pre-trained on 104 languages present relatively lower performance. The KoELECTRA obtains overall the best or second best among six metrics, although this model has a seven times smaller parameter size (14M) than an average of other models (99M). This indicates the effects of the pre-training data source, considering that the KoELECTRA includes the corpus from Namuwiki and Modu that contain modern slang and buzzwords, while other models generally use Korean Wikipedia.

**Evaluation for Multi-labels** Table 6 shows the breakdown F1(micro) for multi-label classification from 1 to 4 labels<sup>7</sup>. A single label task, achieving 0.8553 and 0.8490 from the KR-BERT-c and

KoELECTRA, outperforms other multi-label tasks due to domain similarity. For the multi-label classification, KR-BERT-s achieved the best performance. It uses a sub-character tokenizer that can decompose Hangul(Korean language) syllable characters into sub-characters. Therefore, it provides greater granularity in detecting hate speech words, by identifying the sub-characters from different hate speech categories.<sup>8</sup>

**Evaluation for Label-pairs** Table 7 shows the F1-[macro, micro] scores for curated label pairs based on the proportion in the 2-labels classification. It illustrates that the KR-BERT-s model outperforms in six label pairs. In particular, it is very effective at detecting the *origin and gender* pairs, achieving the highest F1 micro scores of 0.9494 across all label pairs and models. This model uses the sub-character-level tokenizer that can decompose various Korean characters (Hangul syllables) into sub-characters or graphemes to enable handling the bottom consonant (e.g. "gold-digger" [kko#t#baem] 꼬#ㅈ#뱀) or initial consonant (e.g. [k] ㅋ). This approach can detect new slang even if it is only a minor variation from other neutral words.

## 6 Conclusion

We propose K-MHaS, a new large-sized dataset for Korean hate speech detection with a multi-label annotation scheme. We provided extensive baseline experiment results, presenting the usability of a dataset to detect Korean language patterns in hate speech. In future work, the automatic hate speech moderation and counter-speech can be expanded.

<sup>6</sup>The KR-BERT tokenization variants can be found as follows: <https://github.com/snunlp/KR-BERT#tokenization>

<sup>7</sup>Further details are shown in Appendix Table 8.

<sup>8</sup>(e.g.) 개빠르갱이년 = 개 ("dog" - *profanity*) + 빠르갱이 ("communist" - *politics*) + 년 ("bitch" - *gender*)

## Acknowledgements

We thank you the anonymous reviewers for their useful comments. This research is supported by an Australian Government Research Training Program (RTP) Scholarship.

## References

- Tommaso Caselli, Arjan Schelhaas, Marieke Weultjes, Folkert Leistra, Hylke van der Veen, Gerben Timmerman, and Malvina Nissim. 2021. Dalc: the dutch abusive language corpus. In *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH 2021)*, pages 54–66.
- Kimberlé Crenshaw. 1989. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *u. Chi. Legal f.*, page 139.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 11.
- Paula Fortuna, Joao Rocha da Silva, Leo Wanner, Sérgio Nunes, et al. 2019. A hierarchically-labeled portuguese hate speech dataset. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 94–104.
- Paula Fortuna and Sérgio Nunes. 2018. A survey on automatic detection of hate speech in text. *ACM Computing Surveys (CSUR)*, 51(4):1–30.
- Lei Gao and Ruihong Huang. 2017. Detecting online hate speech using context aware models. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP*, pages 260–266.
- Shantanu Godbole and Sunita Sarawagi. 2004. Discriminative methods for multi-labeled classification. In *Pacific-Asia conference on knowledge discovery and data mining*, pages 22–30. Springer.
- Muhammad Okky Ibrohim and Indra Budi. 2019. Multi-label hate speech and abusive language detection in indonesian twitter. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 46–57.
- Myungkoo Kang, Marie-Orange Rivé-Lasan, Wooja Kim, and Philippa Hall. 2020. *Hate speech in Asia and Europe: Beyond hate and fear*. Routledge.
- Sangah Lee, Hansol Jang, Yunmee Baik, Suzi Park, and Hyopil Shin. 2020. Kr-bert: A small-scale korean-specific language model. *ArXiv*, abs/2008.03979.
- Jihyung Moon, Won Ik Cho, and Junbum Lee. 2020. Beep! korean corpus of online news comments for toxic speech detection. In *Proceedings of the Eighth International Workshop on Natural Language Processing for Social Media*, pages 25–31.
- Hamdy Mubarak, Kareem Darwish, and Walid Magdy. 2017. Abusive language detection on arabic social media. In *Proceedings of the first workshop on abusive language online*, pages 52–56.
- Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive language detection in online user content. In *Proceedings of the 25th international conference on world wide web*, pages 145–153.
- 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 EMNLP and the 9th IJCNLP (EMNLP-IJCNLP)*, pages 4675–4684.
- Jangwon Park. 2020. Koelectra: Pretrained electra model for korean. <https://github.com/monologg/KoELECTRA>.
- S. Siegel and N.J. Castellan. 1988. *Nonparametric Statistics for the Behavioral Sciences*. McGraw-Hill.
- SKTBrain. 2019. Kobert: Korean bert pre-trained cased. <https://github.com/monologg/KoELECTRA>.
- Mohammad S Sorower. 2010. A literature survey on algorithms for multi-label learning. *Oregon State University, Corvallis*, 18:1–25.
- Bertie Vidgen and Leon Derczynski. 2020. Directions in abusive language training data, a systematic review: Garbage in, garbage out. *Plos one*, 15(12):e0243300.
- Kunze Wang, Dong Lu, Caren Han, Siqu Long, and Josiah Poon. 2020. Detect all abuse! toward universal abusive language detection models. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6366–6376.
- Zeerak Waseem and Dirk Hovy. 2016. Hateful symbols or hateful people? Predictive features for hate speech detection on Twitter. In *Proceedings of the NAACL student research workshop*, pages 88–93.
- Henry Weld, Guanghao Huang, Jean Lee, Tongshu Zhang, Kunze Wang, Xinghong Guo, Siqu Long, Josiah Poon, and Caren Han. 2021. Conda: a contextual dual-annotated dataset for in-game toxicity understanding and detection. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2406–2416.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers:

State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on EMNLP: System Demonstrations*, pages 38–45. Association for Computational Linguistics.

Ellery Wulczyn, Nithum Thain, and Lucas Dixon. 2017. Ex machina: Personal attacks seen at scale. In *Proceedings of the 26th international conference on world wide web*, pages 1391–1399.

Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. Semeval-2020 task 12: Multilingual offensive language identification in social media (offenseval 2020). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1425–1447.

## A Appendix

**Ethics/Broader Impact Statement** The study follows the ethical policy set out in the ACL code of Ethics<sup>9</sup> and addresses the ethical impact of presenting a new dataset. As described in the data section 2, our annotated dataset is based on the online news comments data publicly available on Kaggle and Github. All annotators were recruited from a crowdsourcing platform. They were informed about hate speech before handling the data. Our instructions allowed them to feel free to leave if they were uncomfortable with the content. With respect to the potential risks, we note that the subjectivity of human annotation would impact on the quality of the dataset.

**The Korean language** The Korean language is morphologically rich and the character structure is different to Latin-based language. A brief components used in the paper as follows:

- **Consonant (자음)** : A consonant is a sound such as ‘p’, ‘f’, ‘n’, or ‘t’ which you pronounce by stopping the air flowing freely through your mouth.
  - initial consonant (초성)
  - bottom consonant (받침)
- **Vowel (모음)** : A vowel is a sound such as the ones represented in writing by the letters ‘a’, ‘e’, ‘i’, ‘o’, and ‘u’, which you pronounce with your mouth open, allowing the air to flow through it.
- **Syllable (음절)** : A syllable is a part of a word that contains a single vowel sound and that is pronounced as a unit. So, for example,

<sup>9</sup><https://www.aclweb.org/portal/content/acl-code-ethics>

‘book’ has one syllable, and ‘reading’ has two syllables.

- Korean romanization : (e.g. [kko#t#baem])
- Character level : (e.g. 꽃#뱀)
- Sub-character level : (e.g. 꼬#≠#뱀)

**Implementation Details** For all baselines, we set the number of epochs as 4 and use a batch size of 32. For other hyper-parameters, we follow the configuration in the official GitHub implementation of the baselines. The source codes or pre-trained models for the baselines are available at the following GitHub addresses: Multilingual BERT<sup>10</sup>, KoELECTRA<sup>11</sup>, KoBERT<sup>12</sup> and KR-BERT<sup>13</sup>.

**Experiments.** A brief of tables displayed in Appendix as follows:

- Table 8: a breakdown of multi-label classification performance from 1 to 4 labels;
- Table 9: overall binary classification performance;
- Table 10: a breakdown of binary classification performance;
- Table 10: F1 score for the top 10 three-label pairs in 3-labels classification;
- Table 12: F1 score for the top 5 four-label pairs in 4-labels classification.

<sup>10</sup><https://github.com/google-research/bert>

<sup>11</sup><https://github.com/monologg/KoELECTRA>

<sup>12</sup><https://github.com/SKTBrain/KoBERT>

<sup>13</sup><https://github.com/snunlp/KR-BERT>

# Labels	Model	F1 (Macro)	F1 (Micro)	F1 (Weighted)	E.M.	AUC	H.L. (↓)
1	BERT	0.6666	0.8190	0.8202	0.7919	0.9011	0.0406
	KoELECTRA	0.6953	<u>0.8490</u>	<u>0.8508</u>	<b>0.8263</b>	<b>0.9213</b>	<u>0.0341</u>
	KoBERT	<u>0.7321</u>	0.8320	0.8370	0.8142	0.9110	0.0379
	KR-BERT(w. char)	<b>0.7336</b>	<b>0.8553</b>	<b>0.8543</b>	<u>0.8239</u>	<u>0.9145</u>	<b>0.0318</b>
	KR-BERT(w. sub)	0.6985	0.8392	0.8419	0.8062	0.9123	0.0360
2	BERT	0.6389	0.8043	0.8174	0.5580	0.8524	0.0788
	KoELECTRA	<u>0.6777</u>	0.8612	0.8700	0.6511	0.8934	0.0577
	KoBERT	<b>0.7249</b>	<b>0.8854</b>	<b>0.8911</b>	<b>0.6794</b>	<b>0.9112</b>	<b>0.0482</b>
	KR-BERT(w. char)	0.6748	0.8405	0.8451	0.5912	0.8735	0.0642
	KR-BERT(w. sub)	0.6718	<u>0.8703</u>	<u>0.8723</u>	<u>0.6535</u>	<u>0.9000</u>	<u>0.0542</u>
3	BERT	0.5784	0.7517	0.7522	0.2448	0.8040	0.1402
	KoELECTRA	0.6146	0.7987	0.7953	0.3310	0.8362	0.1169
	KoBERT	<b>0.6523</b>	<u>0.8290</u>	<u>0.8251</u>	<b>0.3759</b>	<u>0.8589</u>	<u>0.1019</u>
	KR-BERT(w. char)	0.5828	0.7827	0.7732	0.2828	0.8239	0.1230
4	KR-BERT(w. sub)	<u>0.6164</u>	<b>0.8329</b>	<b>0.8263</b>	<u>0.3586</u>	<b>0.8615</b>	<b>0.0996</b>
	BERT	0.4776	0.7093	0.7029	<b>0.1200</b>	0.7610	0.2222
	KoELECTRA	0.4511	0.7044	0.6639	0.0000	0.7680	0.2089
	KoBERT	0.4177	0.6832	0.6460	<u>0.0400</u>	0.7510	0.2267
	KR-BERT(w. char)	<u>0.4837</u>	<u>0.7439</u>	<u>0.7226</u>	<b>0.1200</b>	<u>0.7930</u>	<u>0.1867</u>
KR-BERT(w. sub)	<b>0.5068</b>	<b>0.7771</b>	<b>0.7618</b>	<b>0.1200</b>	<b>0.8120</b>	<b>0.1733</b>	

Table 8: A breakdown of multi-label classification performance from 1 to 4 labels on K-MHaS for the five pre-trained language models at epoch 4 (E.M.:Exact Match, H.L.:Hamming Loss / KR-BERT (w. \*): char = character-level, sub = sub-character-level)

Model	F1 (Macro)	F1 (Micro)	F1 (Weighted)	E.M.	AUC	H.L. (↓)
BERT	0.8495	0.8507	0.8505	0.8507	0.8488	0.1493
KoELECTRA	0.8756	0.8766	0.8765	0.8766	0.8750	0.1234
KoBERT	0.8687	0.8692	0.8693	0.8692	0.8696	0.1308
KR-BERT (w. char)	<u>0.8846</u>	<u>0.8850</u>	<u>0.8851</u>	<u>0.8850</u>	<b>0.8862</b>	<u>0.1150</u>
KR-BERT (w. sub)	<b>0.8869</b>	<b>0.8879</b>	<b>0.8877</b>	<b>0.8879</b>	<u>0.8857</u>	<b>0.1121</b>

Table 9: Overall binary classification performance on the K-MHaS dataset for the five pre-trained language models at epoch 4 (E.M.:Exact Match, H.L.:Hamming Loss / KR-BERT (w. \*): char = character-level, sub = sub-character-level)

Label	Model	F1 (Macro)	F1 (Micro)	F1 (Weighted)	E.M.	H.L. (↓)
Hate Speech	BER	0.4518	0.8243	0.9037T	0.8243	0.1757
	KoELECTRA	0.4606	0.8540	<u>0.9212</u>	0.8540	0.1460
	KoBERT	<u>0.4666</u>	<u>0.8746</u>	<b>0.9331</b>	<u>0.8746</u>	<u>0.1254</u>
	KR-BERT (w. char)	0.4611	0.8558	0.7892	0.8558	0.1442
	KR-BERT (w. sub)	<b>0.4724</b>	<b>0.8953</b>	0.8458	<b>0.8953</b>	<b>0.1047</b>
None	BERT	0.4662	0.8733	<u>0.9323</u>	0.8733	0.1267
	KoELECTRA	<u>0.4726</u>	<u>0.8960</u>	<b>0.9452</b>	<b>0.8960</b>	<u>0.1040</u>
	KoBERT	0.4637	0.8645	0.9273	0.8645	0.1355
	KR-BERT (w. char)	<b>0.4772</b>	<b>0.9126</b>	0.8709	<b>0.9126</b>	<b>0.0874</b>
	KR-BERT (w. sub)	0.4687	0.8821	0.8268	0.8821	0.1179

Table 10: A breakdown of binary classification performance on the K-MHaS dataset for the five pre-trained language models at epoch 4 (E.M.:Exact Match, H.L.:Hamming Loss / KR-BERT (w. \*): char = character-level, sub = sub-character-level, bi = Bidirectional WordPiece tokenizer)



Label Triplets	# triplets	F1 (macro)					F1 (micro)				
		BERT	KoELECTRA	KoBERT	KR-BERT-c	KR-BERT-s	BERT	KoELECTRA	KoBERT	KR-BERT-c	KR-BERT-s
<b>Overall Performance (F1)</b>		0.6912	0.7245	<b>0.7651</b>	<u>0.7444</u>	0.7245	0.8139	<u>0.8493</u>	0.8413	<b>0.8500</b>	0.8445
<i>Origin &amp; Politics &amp; Profanity</i>	41	0.2780	0.2935	<u>0.2954</u>	0.2937	<b>0.3125</b>	0.8224	0.8739	<u>0.8869</u>	0.8739	<b>0.9316</b>
<i>Politics &amp; Profanity &amp; Age</i>	37	0.2781	<u>0.3054</u>	<b>0.3174</b>	0.2981	0.2971	0.8205	<u>0.9126</u>	<b>0.9395</b>	0.8867	0.8824
<i>Physical &amp; Politics &amp; Profanity</i>	32	0.2483	0.2809	<b>0.2960</b>	0.2823	<u>0.2939</u>	0.7296	0.8304	<u>0.8750</u>	0.8421	<b>0.8764</b>
<i>Origin &amp; Profanity &amp; Gender</i>	30	<u>0.2545</u>	0.2314	0.2527	0.2463	<b>0.2886</b>	0.7467	0.7397	0.7368	<u>0.7671</u>	<b>0.8712</b>
<i>Physical &amp; Profanity &amp; Gender</i>	24	0.2151	0.2665	<u>0.2730</u>	0.2459	<b>0.2811</b>	0.6306	0.7869	<u>0.8226</u>	0.7521	<b>0.8413</b>
<i>Origin &amp; Physical &amp; Profanity</i>	14	0.2692	0.2811	<b>0.2873</b>	0.2351	<u>0.2865</u>	0.7532	<u>0.8378</u>	0.8312	0.7273	<b>0.8421</b>
<i>Politics &amp; Age &amp; Gender</i>	14	0.2593	<b>0.2933</b>	<u>0.2830</u>	0.2319	0.2406	0.7606	<b>0.8684</b>	<u>0.8158</u>	0.6970	0.7059
<i>Profanity &amp; Age &amp; Gender</i>	13	0.2686	0.2712	<u>0.2932</u>	0.2593	<b>0.2695</b>	0.8182	0.8060	<b>0.8732</b>	0.8000	<u>0.8358</u>
<i>Origin &amp; Physical &amp; Gender</i>	12	<u>0.2692</u>	0.2327	0.2407	0.2347	<b>0.2703</b>	<u>0.7812</u>	0.6780	0.7458	0.7143	<b>0.8065</b>
<i>Origin &amp; Age &amp; Gender</i>	10	0.2736	<u>0.2781</u>	<b>0.3093</b>	0.2428	0.2411	<u>0.8235</u>	0.8302	<b>0.9123</b>	0.7347	0.7600

Table 11: F1 score for the top 10 three-label pairs on the K-MHaS dataset for the five pre-trained language models at epoch 4 (# total label triplets = 290 / KR-BERT-\*: c = character-level WordPiece tokenizer, s = sub-character-level WordPiece tokenizer)

Label Quadruplets	# quadruplets	F1 (macro)					F1 (micro)				
		BERT	KoELECTRA	KoBERT	KR-BERT-c	KR-BERT-s	BERT	KoELECTRA	KoBERT	KR-BERT-c	KR-BERT-s
<b>Overall Performance (F1)</b>		0.4776	0.4511	0.4177	<u>0.4837</u>	<b>0.5068</b>	0.7093	0.7044	0.6832	<u>0.7439</u>	<b>0.7771</b>
<i>Origin &amp; Profanity &amp; Age &amp; Gender</i>	5	<b>0.3567</b>	0.2950	0.2346	0.2932	<u>0.3086</u>	<b>0.8000</b>	0.6875	0.5806	0.7500	<u>0.7879</u>
<i>Origin &amp; Physical &amp; Profanity &amp; Gender</i>	4	0.2582	0.2804	0.2508	<u>0.3302</u>	<b>0.3757</b>	0.5385	0.7200	0.6400	<u>0.7692</u>	<b>0.8571</b>
<i>Origin &amp; Physical &amp; Politics &amp; Profanity</i>	3	<b>0.4000</b>	<u>0.3333</u>	0.3111	0.3333	<u>0.3667</u>	<b>0.9091</b>	0.8000	0.7368	0.8000	<u>0.8571</u>
<i>Origin &amp; Politics &amp; Profanity &amp; Age</i>	3	<u>0.4222</u>	0.3444	0.3667	<b>0.4444</b>	<b>0.4444</b>	0.8800	0.8000	0.8571	<b>1.0000</b>	<u>0.9600</u>
<i>Origin &amp; Politics &amp; Profanity &amp; Gender</i>	2	0.1852	<b>0.2963</b>	<u>0.2593</u>	<b>0.2963</b>	<b>0.2963</b>	0.5455	<b>0.7692</b>	<u>0.6667</u>	<u>0.6667</u>	<b>0.7692</b>

Table 12: F1 score for the top 5 four-label pairs on the K-MHaS dataset for the five pre-trained language models at epoch 4 (# total label quadruplets = 25 / KR-BERT-\*: c = character-level WordPiece tokenizer, s = sub-character-level WordPiece tokenizer)