

# SemEval-2019 Task 6: Identifying and Categorizing Offensive Language in Social Media (OffensEval)

Marcos Zampieri,<sup>1</sup> Shervin Malmasi,<sup>2</sup> Preslav Nakov,<sup>3</sup>  
Sara Rosenthal,<sup>4</sup> Noura Farra,<sup>5</sup> Ritesh Kumar<sup>6</sup>

<sup>1</sup>University of Wolverhampton, UK, <sup>2</sup>Amazon Research, USA

<sup>3</sup>Qatar Computing Research Institute, HBKU, Qatar

<sup>4</sup>IBM Research, USA, <sup>5</sup>Columbia University, USA, <sup>6</sup>Bhim Rao Ambedkar University, India  
m.zampieri@wlv.ac.uk

## Abstract

We present the results and the main findings of SemEval-2019 Task 6 on Identifying and Categorizing Offensive Language in Social Media (OffensEval). The task was based on a new dataset, the Offensive Language Identification Dataset (OLID), which contains over 14,000 English tweets. It featured three sub-tasks. In sub-task A, the goal was to discriminate between offensive and non-offensive posts. In sub-task B, the focus was on the type of offensive content in the post. Finally, in sub-task C, systems had to detect the target of the offensive posts. OffensEval attracted a large number of participants and it was one of the most popular tasks in SemEval-2019. In total, about 800 teams signed up to participate in the task, and 115 of them submitted results, which we present and analyze in this report.

## 1 Introduction

Recent years have seen the proliferation of offensive language in social media platforms such as Facebook and Twitter. As manual filtering is very time consuming, and as it can cause post-traumatic stress disorder-like symptoms to human annotators, there have been many research efforts aiming at automating the process. The task is usually modeled as a supervised classification problem, where systems are trained on posts annotated with respect to the presence of some form of abusive or offensive content. Examples of offensive content studied in previous work include hate speech (Davidson et al., 2017; Malmasi and Zampieri, 2017, 2018), cyberbullying (Dinakar et al., 2011), and aggression (Kumar et al., 2018). Moreover, given the multitude of terms and definitions used in the literature, some recent studies have investigated the common aspects of different abusive language detection sub-tasks (Waseem et al., 2017; Wiegand et al., 2018).

Interestingly, none of this previous work has studied both the type *and* the target of the offensive language, which is our approach here. Our task, OffensEval<sup>1</sup>, uses the Offensive Language Identification Dataset (OLID)<sup>2</sup> (Zampieri et al., 2019), which we created specifically for this task. OLID is annotated following a hierarchical three-level annotation schema that takes both the target and the type of offensive content into account. Thus, it can relate to phenomena captured by previous datasets such as the one by Davidson et al. (2017). Hate speech, for example, is commonly understood as an insult targeted at a *group*, whereas cyberbullying is typically targeted at an *individual*.

We defined three sub-tasks, corresponding to the three levels in our annotation schema:<sup>3</sup>

**Sub-task A:** Offensive language identification (104 participating teams)

**Sub-task B:** Automatic categorization of offense types (71 participating teams)

**Sub-task C:** Offense target identification (66 participating teams)

The remainder of this paper is organized as follows: Section 2 discusses prior work, including shared tasks related to OffensEval. Section 3 presents the shared task description and the sub-tasks included in OffensEval. Section 4 includes a brief description of OLID based on (Zampieri et al., 2019). Section 5 discusses the participating systems and their results in the shared task. Finally, Section 6 concludes and suggests directions for future work.

<sup>1</sup><http://competitions.codalab.org/competitions/20011>

<sup>2</sup><http://scholar.harvard.edu/malmasi/olid>

<sup>3</sup>A total of 800 teams signed up to participate in the task, but only 115 teams ended up submitting results eventually.

## 2 Related Work

Different abusive and offense language identification problems have been explored in the literature ranging from aggression to cyber bullying, hate speech, toxic comments, and offensive language. Below we discuss each of them briefly.

**Aggression identification:** The TRAC shared task on Aggression Identification (Kumar et al., 2018) provided participants with a dataset containing 15,000 annotated Facebook posts and comments in English and Hindi for training and validation. For testing, two different sets, one from Facebook and one from Twitter, were used. The goal was to discriminate between three classes: non-aggressive, covertly aggressive, and overtly aggressive. The best-performing systems in this competition used deep learning approaches based on convolutional neural networks (CNN), recurrent neural networks, and LSTM (Aroyehun and Gelbukh, 2018; Majumder et al., 2018).

**Bullying detection:** There have been several studies on cyber bullying detection. For example, Xu et al. (2012) used sentiment analysis and topic models to identify relevant topics, and Dadvar et al. (2013) used user-related features such as the frequency of profanity in previous messages.

**Hate speech identification:** This is the most studied abusive language detection task (Kwok and Wang, 2013; Burnap and Williams, 2015; Djuric et al., 2015). More recently, Davidson et al. (2017) presented the hate speech detection dataset with over 24,000 English tweets labeled as non-offensive, hate speech, and profanity.

**Offensive language:** The GermEval<sup>4</sup> (Wiegand et al., 2018) shared task focused on offensive language identification in German tweets. A dataset of over 8,500 annotated tweets was provided for a course-grained binary classification task in which systems were trained to discriminate between offensive and non-offensive tweets. There was also a second task where the offensive class was subdivided into profanity, insult, and abuse. This is similar to our work, but there are three key differences: (i) we have a third level in our hierarchy, (ii) we use different labels in the second level, and (iii) we focus on English.

<sup>4</sup><http://projects.fzai.h-da.de/iggsa/>

**Toxic comments:** The Toxic Comment Classification Challenge<sup>5</sup> was an open competition at Kaggle, which provided participants with comments from Wikipedia organized in six classes: toxic, severe toxic, obscene, threat, insult, identity hate. The dataset was also used outside of the competition (Georgakopoulos et al., 2018), including as additional training material for the aforementioned TRAC shared (Fortuna et al., 2018).

While each of the above tasks tackles a particular type of abuse or offense, there are many commonalities. For example, an insult targeted at an individual is commonly known as cyberbullying and insults targeted at a group are known as hate speech. The hierarchical annotation model proposed in OLID (Zampieri et al., 2019) and used in OffensEval aims to capture this. We hope that the OLID’s dataset would become a useful resource for various offensive language identification tasks.

## 3 Task Description and Evaluation

The training and testing material for OffensEval is the aforementioned Offensive Language Identification Dataset (OLID) dataset, which was built specifically for this task. OLID was annotated using a hierarchical three-level annotation model introduced in Zampieri et al. (2019). Four examples of annotated instances from the dataset are presented in Table 1. We use the annotation of each of the three layers in OLID for a sub-task in OffensEval as described below.

### 3.1 Sub-task A: Offensive language identification

In this sub-task, the goal is to discriminate between offensive and non-offensive posts. Offensive posts include insults, threats, and posts containing any form of untargeted profanity. Each instance is assigned one of the following two labels.

- Not Offensive (NOT): Posts that do not contain offense or profanity;
- Offensive (OFF): We label a post as offensive if it contains any form of non-acceptable language (profanity) or a targeted offense, which can be veiled or direct. This category includes insults, threats, and posts containing profane language or swear words.

<sup>5</sup><http://kaggle.com/c/jigsaw-toxic-comment-classification-challenge>

Tweet	A	B	C
@USER He is so generous with his offers.	NOT	—	—
IM FREEEEEE!!!! WORST EXPERIENCE OF MY FUCKING LIFE	OFF	UNT	—
@USER Fuk this fat cock sucker	OFF	TIN	IND
@USER Figures! What is wrong with these idiots? Thank God for @USER	OFF	TIN	GRP

Table 1: Four tweets from the OLID dataset, with their labels for each level of the annotation model.

### 3.2 Sub-task B: Automatic categorization of offense types

In sub-task B, the goal is to predict the type of offense. Only posts labeled as Offensive (OFF) in sub-task A are included in sub-task B. The two categories in sub-task B are the following:

- Targeted Insult (TIN): Posts containing an insult/threat to an individual, group, or others (see sub-task C below);
- Untargeted (UNT): Posts containing non-targeted profanity and swearing. Posts with general profanity are not targeted, but they contain non-acceptable language.

### 3.3 Sub-task C: Offense target identification

Sub-task C focuses on the target of offenses. Only posts that are either insults or threats (TIN) are considered in this third layer of annotation. The three labels in sub-task C are the following:

- Individual (IND): Posts targeting an individual. It can be a famous person, a named individual or an unnamed participant in the conversation. Insults/threats targeted at individuals are often defined as cyberbullying.
- Group (GRP): The target of these offensive posts is a group of people considered as a unity due to the same ethnicity, gender or sexual orientation, political affiliation, religious belief, or other common characteristic. Many of the insults and threats targeted at a group correspond to what is commonly understood as hate speech.
- Other (OTH): The target of these offensive posts does not belong to any of the previous two categories, e.g., an organization, a situation, an event, or an issue.




Figure 1: Example of a confusion matrix provided in the results package for team NULI, which is the best-performing team for sub-task A.

### 3.4 Task Evaluation

Given the strong imbalance between the number of instances in the different classes across the three tasks, we used the macro-averaged F1-score as the official evaluation measure for all three sub-tasks.

At the end of the competition, we provided the participants with packages containing the results for each of their submissions, including tables and confusion matrices, and tables with the ranks listing all teams who competed in each sub-task. For example, the confusion matrix for the best team in sub-task A is shown in Figure 1.

### 3.5 Participation

The task attracted nearly 800 teams and 115 of them submitted their results. The teams that submitted papers for the SemEval-2019 proceedings are listed in Table 2.<sup>6</sup>

<sup>6</sup>ASE-CSE is for Amrita School of Engineering - CSE.

Team	System Description Paper
Amobee	(Rozenal and Biton, 2019)
ASE-CSE	(Sridharan and T, 2019)
bhanodaig	(Kumar et al., 2019)
BNU-HKBU ...	(Wu et al., 2019)
CAMsterdam	(Aglionby et al., 2019)
CN-HIT-MIT	(Yaojie et al., 2019)
ConvAI	(Pavlopoulos et al., 2019)
DA-LD-Hildesheim	(Modha et al., 2019)
DeepAnalyzer	(la Pea and Rosso, 2019)
Duluth	(Pedersen, 2019)
Emad	(Kebriai et al., 2019)
Embeddia	(Pelicon et al., 2019)
Fermi	(Indurthi et al., 2019)
Ghmerti	(Doostmohammadi et al., 2019)
HAD-Tübingen	(Bansal et al., 2019)
HHU	(Oberstrass et al., 2019)
Hope	(Patras et al., 2019)
INGEOTEC	(Graff et al., 2019)
JCTICOL	(HaCohen-Kerner et al., 2019)
jhan014	(Han et al., 2019)
JTML	(Torres and Vaca, 2019)
JU_ETCE_17_21	(Mukherjee et al., 2019)
KMI_Colig	(Rani and Ojha, 2019)
LaSTUS/TALN	(Altin et al., 2019)
LTL-UDE	(Aggarwal et al., 2019)
MIDAS	(Mahata et al., 2019)
Nikolov-Radivchev	(Nikolov and Radivchev, 2019)
NIT_Agartala_NLP_Team	(Swamy et al., 2019)
NLP	(Kapil et al., 2019)
NLP@UIOWA	(Rusert and Srinivasan, 2019)
NLPR@SRPOL	(Seganti et al., 2019)
nlpUP	(Mitrović et al., 2019)
NULI	(Liu et al., 2019)
SINAI	(Plaza-del Arco et al., 2019)
SSN_NLP	(Thenmozhi et al., 2019)
Stop PropagHate	(Fortuna et al., 2019)
Pardeep	(Singh and Chand, 2019)
techssn	(S et al., 2019)
The Titans	(Garain and Basu, 2019)
TUVD	(Shushkevich et al., 2019)
TüKaSt	(Kannan and Stein, 2019)
UBC-NLP	(Rajendran et al., 2019)
UTFPR	(Paetzold, 2019)
UHH-LT	(Wiedemann et al., 2019)
UM-IU@LING	(Zhu et al., 2019)
USF	(Goel and Sharma, 2019)
UVA Wahoos	(Ramakrishnan et al., 2019)
YNU-HPCC	(Zhou et al., 2019)
YNUWB	(Wang et al., 2019)
Zeyad	(El-Zanaty, 2019)

Table 2: The teams that participated in OffensEval and submitted system description papers.

## 4 Data

Below, we briefly describe OLID, the dataset used for our SemEval-2019 task 6. A detailed description of the data collection process and annotation is presented in Zampieri et al. (2019).

OLID is a large collection of English tweets annotated using a hierarchical three-layer annotation model. It contains 14,100 annotated tweets divided into a training partition of 13,240 tweets and a testing partition of 860 tweets. Additionally, a small trial dataset of 320 tweets was made available before the start of the competition.

A	B	C	Train	Test	Total
OFF	TIN	IND	2,407	100	2,507
OFF	TIN	OTH	395	35	430
OFF	TIN	GRP	1,074	78	1,152
OFF	UNT	—	524	27	551
NOT	—	—	8,840	620	9,460
<b>All</b>			13,240	860	14,100

Table 3: Distribution of label combinations in OLID.

The distribution of the labels in OLID is shown in Table 3. We annotated the dataset using the crowdsourcing platform Figure Eight.<sup>7</sup> We ensured the quality of the annotation by only hiring experienced annotators on the platform and by using test questions to discard annotators who did not achieve a certain threshold. All the tweets were annotated by two people. In case of disagreement, a third annotation was requested, and ultimately we used a majority vote. Examples of tweets from the dataset with their annotation labels are shown in Table 1.

## 5 Results

The models used in the task submissions ranged from traditional machine learning, e.g., SVM and logistic regression, to deep learning, e.g., CNN, RNN, BiLSTM, including attention mechanism, to state-of-the-art deep learning models such as ELMo (Peters et al., 2018) and BERT (Devlin et al.). Figure 2 shows a pie chart indicating the breakdown by model type for all participating systems in sub-task A. Deep learning was clearly the most popular approach, as were also ensemble models. Similar trends were observed for sub-tasks B and C.

<sup>7</sup><https://www.figure-eight.com/>




Figure 2: Pie chart showing the models used in sub-task A. ‘N/A’ indicates that the system did not have a description.

Some teams used additional training data, exploring external datasets such as Hate Speech Tweets (Davidson et al., 2017), toxicity labels (Thain et al., 2017), and TRAC (Kumar et al., 2018). Moreover, seven teams indicated that they used sentiment lexicons or a sentiment analysis model for prediction, and two teams reported the use of offensive word lists. Furthermore, several teams used pre-trained word embeddings from FastText (Bojanowski et al., 2016), from GloVe, including Twitter embeddings from GloVe (Pennington et al., 2014) and from word2vec (Mikolov et al., 2013; Godin et al., 2015).

In addition, several teams used techniques for pre-processing the tweets such as normalizing the tokens, hashtags, URLs, retweets (RT), dates, elongated words (e.g., “Hiiiiii” to “Hi”, partially hidden words (“c00l” to “cool”). Other techniques include converting emojis to text, removing uncommon words, and using Twitter-specific tokenizers, such as the Ark Tokenizer<sup>8</sup> (Gimpel et al., 2011) and the NLTK TweetTokenizer,<sup>9</sup> as well as standard tokenizers (Stanford Core NLP (Manning et al., 2014), and the one from Keras.<sup>10</sup> Approximately a third of the teams indicated that they used one or more of these techniques.

<sup>8</sup><http://www.cs.cmu.edu/~ark/TweetNLP>

<sup>9</sup><http://www.nltk.org/api/nltk.tokenize.html>

<sup>10</sup><http://keras.io/preprocessing/text/>

The results for each of the sub-tasks are shown in Table 4. Due to the large number of submissions, we only show the F1-score for the top-10 teams, followed by result ranges for the rest of the teams. We further include the models and the baselines from (Zampieri et al., 2019): CNN, BiLSTM, and SVM. The baselines are choosing all predictions to be of the same class, e.g., all offensive, and all not offensive for sub-task A. Table 5 shows all the teams that participated in the tasks along with their ranks in each task. These two tables can be used together to find the score/range for a particular team.

Below, we describe the overall results for each sub-task, and we describe the top-3 systems.

## 5.1 Sub-task A

Sub-task A was the most popular sub-task with 104 participating teams. Among the top-10 teams, seven used BERT (Devlin et al.) with variations in the parameters and in the pre-processing steps. The top-performing team, *NULI*, used BERT-base-uncased with default-parameters, but with a max sentence length of 64 and trained for 2 epochs. The 82.9% F1 score of *NULI* is 1.4 points better than the next system, but the difference between the next 5 systems, ranked 2-6, is less than one point: 81.5%-80.6%. The top non-BERT model, *MIDAS*, is ranked sixth. They used an ensemble of CNN and BLSTM+BGRU, together with Twitter word2vec embeddings (Godin et al., 2015) and token/hashtag normalization.

## 5.2 Sub-task B

A total of 76 teams participated in sub-task B, and 71 of them had also participated in sub-task A. In contrast to sub-task A, where BERT clearly dominated, here five of the top-10 teams used an ensemble model. Interestingly, the best team, *jhan014*, which was ranked 76th in sub-task A, used a rule-based approach with a keyword filter based on a Twitter language behavior list, which included strings such as hashtags, signs, etc., achieving an F1-score of 75.5%. The second and the third teams, *Amobee* and *HHU*, used ensembles of deep learning (including BERT) and non-neural machine learning models. The best team from sub-task A also performed well here, ranked 4th (71.6%), thus indicating that overall BERT works well for sub-task B as well.

Sub-task A		Sub-task B		Sub-task C	
Team Ranks	F1 Range	Team Ranks	F1 Range	Team Ranks	F1 Range
1	0.829	1	0.755	1	0.660
2	0.815	2	0.739	2	0.628
3	0.814	3	0.719	3	0.626
4	0.808	4	0.716	4	0.621
5	0.807	5	0.708	5	0.613
6	0.806	6	0.706	6	0.613
7	0.804	7	0.700	7	0.591
8	0.803	8	0.695	8	0.588
9	0.802	9	0.692	9	0.587
<b>CNN</b>	<b>0.800</b>	<b>CNN</b>	<b>0.690</b>	10	0.586
10	0.798	10	0.687	11-14	.571-.580
11-12	.793-.794	11-14	.680-.682	15-18	.560-.569
13-23	.782-.789	15-24	.660-.671	19-23	.547-.557
24-27	.772-.779	<b>BiLSTM</b>	<b>0.660</b>	24-29	.523-.535
28-31	.765-.768	25-29	.640-.655	30-33	.511-.515
32-40	.750-.759	<b>SVM</b>	<b>0.640</b>	34-40	.500-.509
<b>BiLSTM</b>	<b>0.750</b>	30-38	.600-.638	41-47	.480-.490
41-45	.740-.749	39-49	.553-.595	<b>CNN</b>	<b>0.470</b>
46-57	.730-.739	50-62	.500-.546	<b>BiLSTM</b>	<b>0.470</b>
58-63	.721-.729	<b>ALL TIN</b>	<b>0.470</b>	<b>SVM</b>	<b>0.450</b>
64-71	.713-.719	63-74	.418-.486	46-60	.401-.476
72-74	.704-.709	75	0.270	61-65	.249-.340
<b>SVM</b>	<b>0.690</b>	76	0.121	<b>All IND</b>	<b>0.210</b>
75-89	.619-.699	<b>All UNT</b>	<b>0.100</b>	<b>All GRP</b>	<b>0.180</b>
90-96	.500-.590			<b>ALL OTH</b>	<b>0.090</b>
97-103	.422-.492				
<b>All NOT</b>	<b>0.420</b>				
<b>All OFF</b>	<b>0.220</b>				
104	0.171				

Table 4: F1-Macro for the top-10 teams followed by the rest of the teams grouped in ranges for all three sub-tasks. Refer to Table 5 to see the team names associated with each rank. We also include the models (**CNN**, **BiLSTM**, and **SVM**) and the baselines (**All NOT** and **All OFF**) from (Zampieri et al., 2019), shown in bold.

### 5.3 Sub-task C

A total of 66 teams participated in sub-task C, and most of them also participated in sub-tasks A and B. As in sub-task B, ensembles were quite successful and were used by five of the top-10 teams. However, as in sub-task A, the best team, *vradivchev\_anikolov*, used BERT after trying many other deep learning methods. They also used pre-processing and pre-trained word embeddings based on GloVe. The second best team, *NLPR@SRPOL*, used an ensemble of deep learning models such as OpenAI Finetune, LSTM, Transformer, and non-neural machine learning models such as SVM and Random Forest.

### 5.4 Description of the Top Teams

The top-3 teams by average rank for all three sub-tasks were *NLPR@SRPOL*, *NULI*, and *vradivchev\_anikolov*. Below, we provide a brief description of their approaches:

**NLPR@SRPOL** was ranked 8th, 9th, and 2nd on sub-tasks A, B, and C, respectively. They used ensembles of OpenAI GPT, Random Forest, the Transformer, Universal encoder, ELMo, and combined embeddings from fast-Text and custom ones. They trained their models on multiple publicly available offensive datasets, as well as on their own custom dataset annotated by linguists.

Team	Sub-task			Team	Sub-task			Team	Sub-task		
	A	B	C		A	B	C		A	B	C
NULI	<b>1</b>	4	18	resham	40	43	-	kroniker	79	71	-
<u>vrdivchev_anikolov</u>	<u>2</u>	16	<b>1</b>	Xcosmos	41	47	29	aswathyprem	80	-	-
UM-IU@LING	3	76	27	jkolis	42	-	-	DeepAnalyzer	81	38	45
Embeddia	4	18	5	NIT_Agartala_NLP_Team	43	5	38	Code Lyoko	82	-	-
MIDAS	5	8	-	Stop PropagHate	44	-	-	rowantahseen	83	-	-
BNU-HKBU	6	62	39	KVETHZ	45	52	26	ramjib	84	-	-
SentiBERT	7	-	-	christoph.alt	46	14	36	OmerElshrief	85	-	-
<u>NLPR@SRPOL</u>	8	9	<u>2</u>	TECHSSN	47	22	16	desi	86	56	-
YNUWB	9	-	-	USF	48	32	62	Fermi	87	31	3
LTL-UDE	10	-	19	Ziv_Ben_David	49	64	33	mkannan	88	-	-
nlpUP	11	-	-	JCTICOL	50	63	-	mking	89	35	54
ConvAI	12	11	35	TüKaSt	51	23	50	ninab	90	69	-
Vadym	13	10	-	Gal_DD	52	66	25	dianalungu725	91	74	65
UHH-LT	14	21	13	HAD-Tübingen	53	59	61	Halamulki	92	-	-
CAMsterdam	15	19	20	Emad	54	-	-	SSN_NLP	93	65	64
YNU-HPCC	16	-	-	NLP@UIOWA	55	27	37	UTFPR	94	-	-
nishnik	17	-	-	INGEOTEC	56	15	12	rogersdepelle	95	-	-
<u>Amobee</u>	18	<u>2</u>	7	Duluth	57	39	44	Amimul Ihsan	96	-	-
himanisoni	19	46	11	Zeyad	58	34	34	supriyamandal	97	75	-
samsam	20	-	-	ShalomRochman	59	70	58	ramitpahwa	98	-	-
JU_ETCE_17_21	21	50	47	stefaniehegele	60	-	-	ASE - CSE	99	33	32
DA-LD-Hildesheim	22	28	21	NLP-CIC	61	48	46	kripo	100	-	-
YNU-HPCC	23	12	4	Elyash	62	67	40	garain	101	44	63
ChenXiuling	24	-	28	KMLColing	63	45	53	NAYEL	102	-	-
Ghmerti	25	29	-	RUG_OffenseEval	64	-	-	magnito60	103	-	-
safina	26	-	-	jaypee1996	65	41	-	AyushS	104	36	48
Arjun Roy	27	17	-	orabia	66	55	8	UBC_NLP	-	6	9
CN-HIT-MIT	28	30	22	v.gambhir15	67	58	60	bhanodaig	-	57	-
LaSTUS/TALN	29	20	15	kerner-jct.ac.il	68	68	42	Panaetius	-	60	-
HHU	30	3	-	SINAI	69	-	-	eruppert	-	61	-
na14	31	26	10	apalmer	70	13	55	Macporal	-	72	-
NRC	32	37	24	ayman	71	53	57	NoOffense	-	-	6
NLP	33	54	52	Geetika	72	24	-	HHU	-	-	14
JTML	34	-	-	Taha	73	51	59	quanzhi	-	-	17
Arup-Baruah	35	25	31	justhalf	74	-	-	TUVD	-	-	23
UVA_Wahoos	36	42	-	Pardeep	75	7	41	mmfouad	-	-	51
NLP@UniBuc	37	73	49	<b>jhan014</b>	76	<b>1</b>	30	balangheorghe	-	-	56
NTUA-ISLab	38	40	43	liuxy94	77	-	-				
Rohit	39	49	-	ngre1989	78	-	-				

Table 5: All the teams that participated in SemEval-2019 Task 6 with their ranks for each sub-task. The symbol ‘-’ indicates that the team did not participate in some of the subtasks. Please, refer to Table 4 to see the scores based on a team’s rank. The top team for each task is in **bold**, and the second-place team is underlined. Note: ASE - CSE stands for Amrita School of Engineering - CSE, and BNU-HKBU stands for BNU-HKBU UIC NLP Team 2.

**NULI** was ranked 1st, 4th, and 18th on sub-tasks A, B, and C, respectively. They experimented with different models including linear models, LSTM, and pre-trained BERT with fine-tuning on the OLID dataset. Their final submissions for all three subtasks only used BERT, which performed best during development. They also used a number of pre-processing techniques such as hashtag segmentation and emoji substitution.

**vradvichev anikolov** was ranked 2nd, 16th, and 1st on sub-tasks A, B, and C, respectively. They trained a variety of models and combined them in ensembles, but their best submissions for sub-tasks A and C used BERT only, as the other models overfitted. For sub-task B, BERT did not perform as well, and they used soft voting classifiers. In all cases, they used pre-trained GloVe vectors and they also applied techniques to address the class imbalance in the training data.

## 6 Conclusion

We have described SemEval-2019 Task 6 on Identifying and Categorizing Offensive Language in Social Media (OffensEval). The task used OLID (Zampieri et al., 2019), a dataset of English tweets annotated for offensive language use, following a three-level hierarchical schema that considers (i) whether a message is offensive or not (for sub-task A), (ii) what is the type of the offensive message (for sub-task B), and (iii) who is the target of the offensive message (for sub-task C).

Overall, about 800 teams signed up for OffensEval, and 115 of them actually participated in at least one sub-task. The evaluation results have shown that the best systems used ensembles and state-of-the-art deep learning models such as BERT. Overall, both deep learning and traditional machine learning classifiers were widely used. More details about the individual systems can be found in their respective system description papers, which are published in the SemEval-2019 proceedings. A list with references to these publications can be found in Table 2; note, however, that only 50 of the 115 participating teams submitted a system description paper.

As is traditional for SemEval, we have made OLID publicly available to the research community beyond the SemEval competition, hoping to facilitate future research on this important topic.

In fact, the OLID dataset and the SemEval-2019 Task 6 competition setup have already been used in teaching curricula in universities in UK and USA. For example, student competitions based on OffensEval using OLID have been organized as part of Natural Language Processing and Text Analytics courses in two universities in UK: Imperial College London and the University of Leeds. System papers describing some of the students' work are publicly accessible<sup>11</sup> and have also been made available on arXiv.org (Cambray and Podsadowski, 2019; Frisiani et al., 2019; Ong, 2019; Sabora et al., 2019; Puiu and Brabete, 2019; Uglow et al., 2019). Similarly, a number of students in Linguistics and Computer Science at the University of Arizona in USA have been using OLID in their coursework.

In future work, we plan to increase the size of the OLID dataset, while addressing issues such as class imbalance and the small size for the test partition, particularly for sub-tasks B and C. We would also like to expand the dataset and the task to other languages.

## Acknowledgments

We would like to thank the SemEval-2019 organizers for hosting the OffensEval task and for replying promptly to all our inquiries. We further thank the SemEval-2019 anonymous reviewers for the helpful suggestions and for the constructive feedback, which have helped us improve the text of this report.

We especially thank the SemEval-2019 Task 6 participants for their interest in the shared task, for their participation, and for their timely feedback, which have helped us make the shared task a success.

Finally, we would like to thank Lucia Specia from Imperial College London and Eric Atwell from the University of Leeds for hosting the OffensEval competition in their courses. We further thank the students who participated in these student competitions and especially those who wrote papers describing their systems.

The research presented in this paper was partially supported by an ERAS fellowship, which was awarded to Marcos Zampieri by the University of Wolverhampton, UK.

<sup>11</sup><http://scholar.harvard.edu/malmasi/offenseval-student-systems>



## References

- Piush Aggarwal, Tobias Horsmann, Michael Wojatzki, and Torsten Zesch. 2019. LTL-UDE at SemEval-2019 Task 6: BERT and two-vote classification for categorizing offensiveness. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Guy Aglionby, Chris Davis, Pushkar Mishra, Andrew Caines, Helen Yannakoudakis, Marek Rei, Ekaterina Shutova, and Paula Buttery. 2019. CAMsterdam at SemEval-2019 Task 6: Neural and graph-based feature extraction for the identification of offensive tweets. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Lutfiye Seda Mut Altin, Alex Bravo Serrano, and Horacio Saggion. 2019. LaSTUS/TALN at SemEval-2019 Task 6: Identification and categorization of offensive language in social media with attention-based Bi-LSTM model. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Flor Miriam Plaza-del Arco, Dolores Molina-González, Teresa Martín-Valdivia, and Alfonso Ureña-López. 2019. SINAI at SemEval-2019 Task 6: Incorporating lexicon knowledge into SVM learning to identify and categorize offensive language in social media. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Segun Taofeek Aroyehun and Alexander Gelbukh. 2018. Aggression detection in social media: Using deep neural networks, data augmentation, and pseudo labeling. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC)*, pages 90–97.
- Himanshu Bansal, Daniel Nagel, and Anita Soloveva. 2019. HAD-Tübingen at SemEval-2019 Task 6: Deep learning analysis of offensive language on Twitter: Identification and categorization. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. [Enriching word vectors with subword information](#). *CoRR*, abs/1607.04606.
- Pete Burnap and Matthew L Williams. 2015. Cyber hate speech on Twitter: An application of machine classification and statistical modeling for policy and decision making. *Policy & Internet*, 7(2):223–242.
- Aleix Cambay and Norbert Podszadowski. 2019. Bidirectional recurrent models for offensive tweet classification. *arXiv preprint arXiv:1903.08808*.
- Maral Dadvar, Dolf Trieschnigg, Roeland Ordelman, and Franciska de Jong. 2013. Improving cyberbullying detection with user context. In *Advances in Information Retrieval*, pages 693–696. Springer.
- 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 Conference on Weblogs and Social Media (ICWSM)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technology (NAACL-HLT)*.
- Karthik Dinakar, Roi Reichart, and Henry Lieberman. 2011. Modeling the detection of textual cyberbullying. In *The Social Mobile Web*, pages 11–17.
- Nemanja Djuric, Jing Zhou, Robin Morris, Mihajlo Grbovic, Vladan Radosavljevic, and Narayan Bhamidipati. 2015. Hate speech detection with comment embeddings. In *Proceedings of the Web Conference (WWW)*.
- Ehsan Doostmohammadi, Hossein Sameti, and Ali Safar. 2019. Ghmert at SemEval-2019 Task 6: A deep word- and character-based approach to offensive language identification. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Zeyad El-Zanaty. 2019. Zeyad at SemEval-2019 Task 6: That’s offensive! An all-out search for an ensemble to identify and categorize offense in tweets. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Paula Fortuna, José Ferreira, Luiz Pires, Guilherme Routar, and Sérgio Nunes. 2018. Merging datasets for aggressive text identification. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC)*, pages 128–139.
- Paula Fortuna, Juan Soler-Company, and Nunes Srgio. 2019. Stop PropagHate at SemEval-2019 Tasks 5 and 6: Are abusive language classification results reproducible? In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Nicolò Frisiani, Alexis Laignelet, and Batuhan Güler. 2019. Combination of multiple deep learning architectures for offensive language detection in tweets. *arXiv preprint arXiv:1903.08734*.
- Avishek Garain and Arpan Basu. 2019. The Titans at SemEval-2019 Task 6: Hate speech and target detection. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Spiros V Georgakopoulos, Sotiris K Tasoulis, Aris-tidis G Vrahatis, and Vassilis P Plagianakos. 2018. Convolutional neural networks for toxic comment classification. *arXiv preprint arXiv:1802.09957*.

- Kevin Gimpel, Nathan Schneider, Brendan O'Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. 2011. Part-of-speech tagging for Twitter: annotation, features, and experiments. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Frédéric Godin, Baptist Vandersmissen, Wesley De Neve, and Rik Van de Walle. 2015. Multimedia Lab @ACL WNUT NER Shared Task: Named entity recognition for Twitter microposts using distributed word representations. In *Proceedings of the Workshop on Noisy User-generated Text*.
- Bharti Goel and Ravi Sharma. 2019. USF at SemEval-2019 Task 6: Offensive language detection using LSTM with word embeddings. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Mario Graff, Sabino Miranda-Jiménez, Eric S. Tellez, and Daniela Moctezuma. 2019. INGEOTEC at SemEval-2019 Task 5 and Task 6: A genetic programming approach for text classification. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Yaakov HaCohen-Kerner, Ziv Ben-David, Gal Didi, Eli Cahn, Shalom Rochman, and Elyashiv Shayovitz. 2019. JCTICOL at SemEval-2019 Task 6: Classifying offensive language in social media using deep learning methods, word/character n-gram features, and preprocessing methods. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Jiahui Han, Xinyu Liu, and Shengtian Wu. 2019. jhan014 at SemEval-2019 Task 6: Identifying and categorizing offensive language in social media. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Vijayaradhil Indurthi, Bakhtiyar Syed, Manish Shrivastava, Manish Gupta, and Vasudeva Varma. 2019. Fermi at SemEval-2019 Task 6: Identifying and categorizing offensive language in social media using sentence embeddings. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Madeeswaran Kannan and Lukas Stein. 2019. TüKaSt at SemEval-2019 Task 6: something old, something neu(ral): Traditional and neural approaches to offensive text classification. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Prashant Kapil, Asif Ekbal, and Dipankar Das. 2019. NLP at SemEval-2019 Task 6: Detecting offensive language using neural networks. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Emad Kebriaei, Samaneh Karimi, Nazanin Sabri, and Azadeh Shakery. 2019. Emad at SemEval-2019 Task 6: Offensive language identification using traditional machine learning and deep learning approaches. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Ritesh Kumar, Guggilla Bhanodai, Rajendra Pamula, and Chennuru Maheshwar Reddy. 2019. bhanodaig at SemEval-2019 Task 6: Categorizing offensive language in social media. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Ritesh Kumar, Atul Kr Ojha, Shervin Malmasi, and Marcos Zampieri. 2018. Benchmarking aggression identification in social media. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC)*.
- Irene Kwok and Yuzhou Wang. 2013. Locate the hate: Detecting tweets against blacks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Ping Liu, Wen Li, and Liang Zou. 2019. NULI at SemEval-2019 Task 6: Transfer learning for offensive language detection using bidirectional transformers. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Debanjan Mahata, Haimin Zhang, Karan Uppal, Yaman Kumar, Rajiv Ratn Shah, Simra Shahid, Laiba Mehnaz, and Sarthak Anand. 2019. MIDAS at SemEval-2019 Task 6: Identifying offensive posts and targeted offense from Twitter. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Prasenjit Majumder, Thomas Mandl, et al. 2018. Filtering aggression from the multilingual social media feed. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC)*, pages 199–207.
- Shervin Malmasi and Marcos Zampieri. 2017. Detecting Hate Speech in Social Media. In *Proceedings of the Conference on Recent Advances in Natural Language Processing (RANLP)*.
- Shervin Malmasi and Marcos Zampieri. 2018. Challenges in Discriminating Profanity from Hate Speech. *Journal of Experimental & Theoretical Artificial Intelligence*, 30:1 – 16.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of the International Conference on Neural Information Processing Systems (NIPS)*.

- Jelena Mitrović, Bastian Birkeneder, and Michael Granitzer. 2019. nlpUP at SemEval-2019 Task 6: a deep neural language model for offensive language detection. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Sandip Modha, Prasenjit Majumder, and Daksh Patel. 2019. DA-LD-Hildesheim at SemEval-2019 Task 6: Tracking offensive content with deep learning model using shallow representation. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Preeti Mukherjee, Mainak Pal, Somnath Banerjee, and Sudip Kumar Naskar. 2019. JU\_ETCE\_17\_21 at SemEval-2019 Task 6: Efficient machine learning and neural network approaches for identifying and categorizing offensive language in tweets. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Alex Nikolov and Victor Radivchev. 2019. Nikolov-Radivchev at SemEval-2019 Task 6: Offensive tweet classification with BERT and ensembles. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Alexander Oberstrass, Julia Romberg, Anke Stoll, and Stefan Conrad. 2019. HHU at SemEval-2019 Task 6: Context does matter - tackling offensive language identification and categorization with ELMo. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Ryan Ong. 2019. Offensive language analysis using deep learning architecture. *arXiv preprint arXiv:1903.05280*.
- Gustavo Henrique Paetzold. 2019. UTFPR at SemEval-2019 Task 6: Relying on compositionality to find offense. In *Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval)*.
- Gabriel Florentin Patras, Diana Florina Lungu, Daniela Gifu, and Diana Trandabat. 2019. Hope at SemEval-2019 Task 6: Mining social media language to discover offensive language. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- John Pavlopoulos, Nithum Thain, Lucas Dixon, and Ion Androutsopoulos. 2019. ConvAI at SemEval-2019 Task 6: Offensive language identification and categorization with perspective and BERT. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Ted Pedersen. 2019. Duluth at SemEval-2019 Task 6: Lexical approaches to identify and categorize offensive tweets. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Andraž Pelicon, Matej Martinc, and Petra Kralj Novak. 2019. Embeddia at SemEval-2019 Task 6: Detecting hate with neural network and transfer learning approaches. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technology (NAACL-HLT)*.
- Gretel Liz De la Pea and Paolo Rosso. 2019. DeepAnalyzer at SemEval-2019 Task 6: A deep learning-based ensemble method for identifying offensive tweets. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Andrei-Bogdan Puiu and Andrei-Octavian Brabete. 2019. Towards NLP with deep learning: Convolutional neural networks and recurrent neural networks for offensive language identification in social media. *arXiv preprint arXiv:1903.00665*.
- Arun Rajendran, Chiyu Zhang, and Muhammad Abdul-Mageed. 2019. UBC-NLP at SemEval-2019 Task 6: Ensemble learning of offensive content with enhanced training data. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Murugesan Ramakrishnan, Wlodek Zadrozny, and Narges Tabari. 2019. UVA Wahooos at SemEval-2019 Task 6: Hate speech identification using ensemble machine learning. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Priya Rani and Atul Kr. Ojha. 2019. KMIColing at SemEval-2019 Task 6: Exploring n-grams for offensive language detection. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Alon Rozental and Dadi Biton. 2019. Amobee at SemEval-2019 Tasks 5 and 6: Multiple choice CNN over contextual embedding. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Jonathan Rusert and Padmini Srinivasan. 2019. NLP@UIOWA at SemEval-2019 Task 6: Classifying the classless using multi-windowed CNNs. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Angel Deborah S, Rajalakshmi S, Logesh B, Harshini S, Geetika B, Dyaneswaran S, S Milton Rajendram, and Mirnalinee T T. 2019. TECHSSN at SemEval-2019 Task 6: Identifying and categorizing offensive language in tweets using deep neural networks. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.

- Silvia Sapora, Bogdan Lazarescu, and Christo Lolov. 2019. Absit invidia verbo: Comparing deep learning methods for offensive language. *arXiv preprint arXiv:1903.05929*.
- Alessandro Seganti, Helena Sobol, Iryna Orlova, Hannam Kim, Jakub Staniszewski, Tymoteusz Krumholz, and Krystian Koziel. 2019. NLPR@SRPOL at SemEval-2019 Task 6 and Task 5: Linguistically enhanced deep learning offensive sentence classifier. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Elena Shushkevich, John Cardiff, and Paolo Rosso. 2019. TUV D team at SemEval-2019 Task 6: Offense target identification. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Pardeep Singh and Satish Chand. 2019. Pardeep at SemEval-2019 Task 6: Identifying and categorizing offensive language in social media using deep learning. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Murali Sridharan and Swapna T. 2019. Amrita School of Engineering - CSE at SemEval-2019 Task 6: Manipulating attention with temporal convolutional neural network for offense identification and classification. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Steve Durairaj Swamy, Anupam Jamatia, Björn Gambäck, and Amitava Das. 2019. NIT\_Agartala\_NLP\_Team at SemEval-2019 Task 6: An ensemble approach to identifying and categorizing offensive language in Twitter social media corpora. In *Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval)*.
- Nithum Thain, Lucas Dixon, and Ellery Wulczyn. 2017. [Wikipedia Talk Labels: Toxicity](#).
- D Thenmozhi, Senthil Kumar B, Srinethe Sharavanan, and Aravindan Chandrabose. 2019. SSN\_NLP at SemEval-2019 Task 6: Offensive language identification in social media using machine learning and deep learning approaches. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Johnny Torres and Carmen Vaca. 2019. JTML at SemEval-2019 Task 6: Offensive tweets identification using convolutional neural networks. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Harrison Uglow, Martin Zlocha, and Szymon Zmislony. 2019. An exploration of state-of-the-art methods for offensive language detection. *arXiv preprint arXiv:1903.07445*.
- Bin Wang, Xiaobing Zhou, and Xuejie Zhang. 2019. YNUWB at SemEval-2019 Task 6: K-max pooling cnn with average meta-embedding for identifying offensive language. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Zeerak Waseem, Thomas Davidson, Dana Warmusley, and Ingmar Weber. 2017. Understanding abuse: A typology of abusive language detection subtasks. *arXiv preprint arXiv:1705.09899*.
- Gregor Wiedemann, Eugen Ruppert, and Chris Bieermann. 2019. UHH-LT at SemEval-2019 Task 6: Supervised vs. unsupervised transfer learning for offensive language detection. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Michael Wiegand, Melanie Siegel, and Josef Ruppenhofer. 2018. Overview of the GermEval 2018 shared task on the identification of offensive language. In *Proceedings of the GermEval 2018 Workshop (GermEval)*.
- Zhenghao Wu, Hao Zheng, Jianming Wang, Weifeng Su, and Jefferson Fong. 2019. BNU-HKBU UIC NLP Team 2 at SemEval-2019 Task 6: Detecting offensive language using BERT model. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Jun-Ming Xu, Kwang-Sung Jun, Xiaojin Zhu, and Amy Bellmore. 2012. Learning from bullying traces in social media. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technology (NAACL-HLT)*.
- Zhang Yaojie, Xu Bing, and Zhao Tiejun. 2019. CN-HIT-MI.T at SemEval-2019 Task6: Offensive language identification based on BiLSTM with double attention. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. Predicting the type and target of offensive posts in social media. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technology (NAACL-HLT)*.
- Chengjin Zhou, Jin Wang, and Xuejie Zhang. 2019. YNU-HPCC at SemEval-2019 Task 6: Identifying and categorising offensive language on Twitter. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Jian Zhu, Zuoyu Tian, and Sandra Kübler. 2019. UM-IU@LING at SemEval-2019 Task 6: Identifying offensive tweets using BERT and SVMs. In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.