

Leveraging World Knowledge in Implicit Hate Speech Detection

Jessica Lin

Department of Linguistics
Georgetown University
y11290@georgetown.edu

Abstract

Warning: *This paper contains content that may be offensive or disturbing.*

While much attention has been paid to identifying explicit hate speech, implicit hateful expressions that are disguised in coded or indirect language are pervasive and remain a major challenge for existing hate speech detection systems. This paper presents the first attempt to apply Entity Linking (EL) techniques to both explicit and implicit hate speech detection, where we show that such real world knowledge about entity mentions in a text does help models better detect hate speech, and the benefit of adding it into the model is more pronounced when explicit entity triggers (e.g., rally, KKK) are present. We also discuss cases where real world knowledge does not add value to hate speech detection, which provides more insights into understanding and modeling the subtleties of hate speech.

1 Introduction

Hate speech on social media facilitates the spread of violence in the real world. For this reason, the detection of hatred content online increasingly gains importance. However, most work in hate speech detection has focused on explicit or overt hate speech, failing to capture the implicit hateful messages in coded or indirect language (e.g., sarcasm or metaphor) that disparage a protected group or individual, or to convey prejudicial and harmful views about them (Waseem et al., 2017). Examples (1) and (2) from ElSherief et al. (2021) show the two types of hate speech, explicit vs. implicit:

- (1) *#jews & n*ggers destroy and pervert everything they touch #jewfail #n*ggerfail* (explicit hate speech)
- (2) *don't worry, charlottesville was just the beginning. we're growing extremely fast* (implicit hate speech, implied statement: larger white supremacist events will happen)

As shown in (1) and (2), explicit hate speech is direct and uses specific keywords while implicit hate speech does not contain explicit hateful lexicon or phrases and often uses coded or indirect languages to disguise the malicious intent (ElSherief et al., 2021).

Modeling implicit sentiment in hate speech is still in its infancy, and the capacity to acquire background knowledge enhances the correct detection of hate speech by machines (Kiritchenko et al., 2021; Li et al., 2021). To be able to understand the implied statement of hate speech, machine learning systems need extratextual information that provides world knowledge associated with natural language concepts. For example, it would be impossible for a reader who does not know what happened in Charlottesville to understand the implicit hateful message in (2). The reader wouldn't be able to understand the implied message "larger white supremacist events will happen" without knowing that Charlottesville is a metonym for a white supremacist rally that took place in Charlottesville, Virginia in August, 2017. Conversely, background knowledge that reasons about entity mentions in a text could add value to the detection of implicit hate speech. Incorporating such knowledge ideally should make it easier for the learning model to detect hate speech where it is not apparent from the text.

With this motivation, this study applies Entity Linking (EL) to identify entities in tweets, link them to an external knowledge base (KB; Wikipedia in this study), and acquire their Wikipedia descriptions that would be encoded with Sentence BERT (Reimers and Gurevych, 2019) for representation. Our proposed model incorporates such knowledge representation into identifying both explicit and implicit hate speech in investigating the effectiveness of real world knowledge.

Overall, this study makes the following contributions: (i) To the best of our knowledge, this work

is the first attempt to leverage EL techniques in tackling the problem of implicit hate speech detection. (ii) To evaluate the effectiveness of real world knowledge in both explicit and implicit hate speech detection, where we investigate how incorporating Wikipedia descriptions of linked entities into the model affects performance.

2 Related Work

Identifying hate speech has been a topic of immense interest in recent years, and a number of studies have approached this problem in different ways.

Early work on hate speech detection has focused on explicitly abusive text using keyword-based methods that rely on lexical features (Waseem and Hovy, 2016; Davidson et al., 2017), while more recent studies have highlighted the linguistic nuance and diversity of the implicit hate expressions, which includes stereotypes (Sap et al., 2019), indirect sarcasm, humor, and metaphor (Founta et al., 2018) that cannot be captured by keyword-based systems. Implicit hate expressions are no less harmful than explicit ones and make up a large portion of false negatives errors (Basile et al., 2019; Mozafari et al., 2020). Systems that rely on explicit hateful lexicon or phrases are unable to capture underlying hateful intent like humans. Up until now, predicting implicit hate or abuse remains a major challenge for machine systems. Existing solutions for identifying implicit cases of hate speech involve taking context into account. For example, Gao and Huang (2017) included original news articles as the context of the hateful comments. Other studies have built datasets with “implicit” labels or annotations (Caselli et al., 2020; ElSherief et al., 2021; Sap et al., 2019). This is crucial not only for evaluation but also for training, as systems that are not trained on implicit hate would not go beyond explicit features and are thus far from being applicable in the real world as a moderation tool.

Recently, an emerging line of research has started to explore the idea of incorporating real world knowledge in a related task, sarcasm detection, but not for hate speech detection task. This line of research (Chowdhury and Chaturvedi, 2021; Li et al., 2021) hypothesizes that infusing real world knowledge such as commonsense knowledge in sarcasm detection ideally should make the learning model easier to detect sarcasm where it is not apparent from the text. Li et al. (2021)

proposed a novel architecture to integrate knowledge into learning model. For knowledge representation, they applied the pre-trained COMET model (COMMONsense Transformers, Bosselut et al. (2019)) to generate relevant commonsense knowledge from sarcastic instances and use it as input to the proposed model in investigating how commonsense knowledge influences performance. Similarly, Chowdhury and Chaturvedi (2021) leveraged COMET to infuse commonsense knowledge in their graph convolution-based model, in which a graph is formed with edges between the input sentence and COMET sequences. The node representations of the graph are then passed through a fully-connected neural network to generate the output.

The results of the effectiveness of commonsense knowledge are still inconclusive. Li et al. (2021) found that integrating commonsense knowledge information contributes to sarcasm detection, yet it only plays a supporting role as models using only knowledge information do not perform satisfactorily. Interestingly, Chowdhury and Chaturvedi (2021) found an opposite result on the role of commonsense knowledge in sarcasm detection, showing that COMET infused model performs at par with the baseline. In many cases, the model is more reliant on the input sentence and less on the COMET sequences for making the prediction. Their in-depth error analysis shows that commonsense is most effective in identifying sarcasm with polarity contrast but fails to explain non-sarcastic samples or other types of sarcasm effectively. The study suggests that exploring the utility of other forms of external knowledge such as factual world knowledge for sarcasm detection would be a promising line of inquiry. Inspired by previous research, the present study assumes that real world knowledge would be beneficial to implicit hate speech detection in that such knowledge allows a fully understanding of the background knowledge required for deciphering the hateful intent of the text.

3 Methods

3.1 Dataset

In this work, our proposed model is evaluated on the *Latent Hatred Dataset* (ElSherief et al., 2021). The dataset used in this study contains 12,143 tweets from the most prominent extremist groups in the United States, where 5,791 of these tweets are

implicit hate speech and 567 of them are explicit hate speech. Implicit hate tweets are categorized into six classes using the taxonomy shown in Table 1.

Each of the 5,791 implicit hate tweets also has free-text annotations for the target demographic group and an implied statement to describe the underlying message. Implied statements are generated by human annotators with the format $\langle target \rangle \{do, are, commit\} \langle predicate \rangle$, where $\langle target \rangle$ might be phrases like *immigrants, minorities*. For example, the implicit hate tweet *this selfie is so white, i love it* has the implied statement “Minorities are less than whites”.

3.2 Models

In this paper, two classification tasks are conducted. (1) a binary classification task on distinguishing hate speech from non-hate speech, and (2) a 6-way classification task on categorizing implicit hate speech classes (see Table 1).

For both tasks, a Multi-layer Perceptron (MLP) model with Sentence BERT (Reimers and Gurevych, 2019) embeddings is used. First, we pre-processed all tweets and background knowledge descriptions (remove stop words and reserved words such as RT, FAV, via, etc. while keeping all the hashtags and links that may contain useful messages). Next, they are concatenated before being encoded with Sentence BERT using the pre-trained model `bert-base-nli-mean-tokens`, where it maps tweets and background knowledge descriptions to a 768-dimensional dense vector space. The MLP model is evaluated with two feature sets: Sentence BERT encoded textual embedding alone (baseline, tweet only) and the combination of textual embedding and background knowledge (baseline+BK).

3.3 Background Knowledge Extraction and Representation

To incorporate background knowledge, entity linking is applied to associate mentions with their referent entities. First, mentions in each tweet are identified and linked to entities in the KB using Radboud Entity Linker (REL, van Hulst et al. (2020)), an end-to-end entity linker that identifies mentions of specific entities in text and links them to pertinent Wikipedia page titles. REL is chosen because it has state-of-the-art performance and is trained on a recent Wikipedia dump (2019-07). It provides

a web API ¹ in which given an input text it returns a list of mentions with the linked entities and the confidence score of mention detection and entity disambiguation. In order to refine the entity linking results (see Table 2 for an example), we tested different thresholds of confidence score and decided to remove entities with a low confidence score of mention detection (MD score) (<0.4) and a low confidence score of entity disambiguation (ED score) (<0.2). We find that this refinement strategy helps us strike a balance between precision and recall in that it matches as many mentions as possible (retain mentions that have an MD score > 0.4) while maintaining the accuracy of the result at the same time (remove entities that have an ED score < 0.2).

After retrieving all the Wikipedia page titles of the entities in the input text, we use Wikipedia API ² to extract the summary of the corresponding Wikipedia page (referred to as *entity abstract* or *Wikipedia description* afterward). To keep the entity abstract from being too long, we print only two sentences of each abstract by setting the sentences argument to 2. An example of entity abstract is as follows: *David Ernest Duke (born July 1, 1950) is an American white supremacist, antisemitic conspiracy theorist, far-right politician, convicted felon, and former Grand Wizard of the Knights of the Ku Klux Klan. From 1989 to 1992, he was a member of the Louisiana House of Representatives for the Republican Party.*

Finally, the entity abstract for each tweet is concatenated with the tweet and encoded with 768-dimensional Sentence BERT embedding using `bert-base-nli-mean-tokens` model.

4 Experiments

4.1 Experimental Setup

For each of the two classification tasks, the model is trained and evaluated on two feature sets, which are baseline feature set (tweet text only) with and without Wikipedia descriptions. For both binary and 6-way classification task, a MLP is implemented in `sklearn` with three hidden layers of dimension 512, learning rate 0.001 and the number of epochs 500. The optimizer is set to Adam.

¹<https://rel.cs.ru.nl/api>

²<https://pypi.org/project/wikipedia/>

Class (distribution)	Explanation	Example
Grievance (24.2%)	Frustration over a minority group’s perceived privilege	<i>Black lives matter and white lives don’t? Sounds racist.</i>
Incitement (20%)	Implicitly promoting known hate groups and ideologies (e.g., by flaunting in-group power)	<i>Hitler was Germany –Germans shall rise again!</i>
Inferiority (13.6%)	Implying some group or person is of lesser value than another	<i>It’s not a coincidence the best places to live are majority white.</i>
Irony (12.6%)	Using sarcasm, humor, and satire to demean someone	<i>Horrors... Disney will be forced into hiring Americans. (Discredit Disney for allegedly hiring only non-whites)</i>
Stereotypes (17.9%)	Associating a group with negative attribute using euphemisms, circumlocution, or metaphorical language	<i>Can someone tell the black people in Chicago to stop killing one another before it becomes Detroit?</i>
Threats (10.5%)	Making an indirect commitment to attack someone’s body, well-being, reputation, liberty, etc. Focus on threats related to implicit violation of rights and freedoms, removal of opportunities, and more subtle forms of intimidation	<i>All immigration of non-whites should be ended. (Non-white immigrants should stay in their country: subtle forms of intimidation)</i>

Table 1: Implicit hate classes and examples in *Latent Hatred Dataset*.

Tweet	Before Refinement Strategy	After Refinement Strategy (Removing entities with a MD score <0.4 and a ED score <0.2)
tune in today’s jmt for my interview w /robert spencer on “the complete infidel’s guide to iran!”	tune in today’s [jmt]Jedi_Mind_Tricks (an American hip hop group) (ED score: 0.78, MD score: 0.36) for my interview w /[robert spencer] Robert B. Spencer (American author and blogger, opponent of Islam) (ED score: 0.38, MD score: 0.96) on “the complete infidel’s guide to [iran] Iran” (ED score: 0.51, MD score: 0.99)	tune in today’s jmt for my interview w /[robert spencer] Robert B. Spencer (American author and blogger, opponent of Islam) (ED score: 0.38, MD score: 0.96) on “the complete infidel’s guide to [iran] Iran” (ED score: 0.51, MD score: 0.99)

Table 2: An Entity linking example from our dataset.

4.2 Classification Results

In explicit hate speech classification shown in Table 3, the background knowledge provided by Wikipedia significantly improves the model by increasing 10% in precision, recall, and F1 score after incorporating background knowledge into the model. While the MLP model with baseline feature set achieves a competitive result with 65% on F1 score, the background knowledge incorporated model achieves better scores (75%) than the one with baseline feature set, demonstrating that real world knowledge is helpful for capturing real hate speech.

For additional comparisons ³, our background knowledge incorporated model achieves a significantly better precision score (75% vs. 68%) than

³We notice that there are works (Pal et al., 2022) that also improve performance on the *Latent Hatred Dataset*, but we compare our results against the results from ElSherief et al. (2021), which serves as a benchmark for modeling implicit hate speech using knowledge-based features as well.

the Wikidata Knowledge Graph (Vrandečić and Krötzsch, 2014) infused model proposed in ElSherief et al. (2021), which was trained on the same *Latent Hatred Dataset*. The remarkably higher precision score suggests that Wikipedia description of linked entities is doing a better job in preventing false positives than Wikidata Knowledge Graph method; however, a more detailed analysis comparing the effectiveness of different external knowledge (e.g., knowledge graphs, commonsense knowledge) for hate speech detection is needed.

As shown in Table 3, real world knowledge enhances the correct detection of explicit hate speech. However, Table 4 shows that integrating real world knowledge does not seem to improve the model, and even hurts model’s performance in implicit hate speech type classification. Significant degradation in precision, recall, and F1 score are observed (e.g., recall drops 12%), which suggests that knowledge about the involved entities is not sufficient for predicting implicit hate speech types.

Models	P	R	F	Acc
Majority baseline	52 % \pm 1.3 %			
MLP (baseline)	65 % \pm 1.5 %	65 % \pm 1.5 %	65 % \pm 1.5 %	65 % \pm 1.5 %
MLP (baseline+BK)	75 % \pm 1.4 %	75 % \pm 1.4 %	75 % \pm 1.4 %	75 % \pm 1.4 %
Knowledge infused model (ElSherief et al., 2021)	68%	72%	70%	77%

Table 3: Classification performance on explicit hate speech classification. Performance metrics are all macro average scores. Majority baseline always returns the positive (hate) label. **Bold** face indicates best performance.

Table 5 further shows that among the six implicit hate classes, *irony* is the hardest for the model to detect, which aligns with the result found in ElSherief et al. (2021). This is reasonable, as irony normally requires further understanding beyond knowledge about the involved entities (e.g., semantic inference or pragmatic understanding). Therefore, our background knowledge incorporated model fails at capturing this type of implicit hate. On the other hand, *white grievance* is the easiest to detect for our model. A detailed examination of our data shows that compared to other types of implicit hate posts, *white grievance* tweets in our dataset contain relatively more explicit hate triggers (e.g., rally, KKK), which is found to be useful for the model. Further discussion on explaining our model’s predictions on implicit hate could be found in Section 5.

5 Analysis

This section investigates whether the model is reliant on Wikipedia descriptions while making decisions. We leverage LIME (locally interpretable model-agnostic explanations) algorithms (Ribeiro et al., 2016) to explain our model’s predictions on both explicit and implicit hate statements through random examples picked from our dataset.

5.1 Efficacy of World Knowledge

Table 6 shows cases where Wikipedia knowledge is helpful. For each example, the colored text spans represent the words highly-weighted by the model. We find that Wikipedia knowledge is particularly useful when hatefulness in a tweet is conveyed through certain hate “triggers”. These triggers by themselves are not toxic but are relevant to the hatefulness in a tweet. Since implicit hate does not contain explicit hate lexicon or phrases, the model rather relies on these triggers to help them make the right predictions. As shown in Example C of Table 6, the Wikipedia description of *Charlottesville* helps the model correctly predicts the tweet as an incitement tweet by relying on entity triggers such

as *rally*. The word by itself is not a hate lexicon but indicates a high probability of a tweet that incites violence. Similarly, the Wikipedia description of *David Duke* in Example B of Table 6 is helpful for the model in that it explains *David Duke* is the former head of *Ku Klux Klan*, which by itself does not convey toxicity but is indicative of a *white grievance* tweet.

Table 6 further shows that although entity triggers provided by Wikipedia description contribute to the detection of hate speech to a certain extent, it only plays a supporting role. Some of the words in the tweet are already a strong signal of the hatefulness of the tweet, as shown in the highlighted words in the table. For instance, in Example B of Table 6, both the hashtag *#makeamericagreatagain* and *#votetrump* in the post reveal that the author might be a supporter of Donald Trump, which is said to have a symbiotic relationship with white nationalism, white supremacy, and white power ideologies that correspond to the *white grievance* implicit hate type in our dataset.

5.2 Error Analysis

To further understand the role of world knowledge in identifying hatefulness, we randomly pick out incorrect predictions made by our model and manually correct some of the entity linking errors to see if this “post-processing” helps avert classification errors.

As shown in Example A of Table 7, our entity linker misses the mention *bernie bros*. Instead, *white males* in the tweet is identified and linked to Wikipedia. The Wikipedia description of *white males* does not add value to the detection of hatefulness. Words with the highest coefficients such as *skin* and *African* are neutral and are not associated with the hateful content of the tweet. After post-processing the entity linking result, our model correctly predicts the tweet as a hate post by leveraging the world knowledge provided by Wikipedia. The Wikipedia description explains that *bernie bros*

Feature sets	P	R	F	Acc
Dummy Classifier	19 % ± 1.0 %			
Baseline (tweet only)	52 % ± 1.3 %	52 % ± 1.3 %	52 % ± 1.3 %	54% ± 1.3 %
Baseline + BK	42 % ± 1.3 %	40 % ± 1.3 %	41 % ± 1.3 %	44 % ± 1.3 %

Table 4: Classification performance on 6-way implicit hate speech classification. Performance metrics are all macro average scores. Dummy classifier generates random predictions by respecting the training set class distribution. **Bold** face indicates best performance.

Feature sets	incitement	inferiority	irony	stereotypical	threatening	white grievance
MLP (baseline)	52%	51%	33%	55%	56%	62%
MLP (baseline+BK)	45%	40%	20%	44%	43%	52%

Table 5: Classification performance on 6-way implicit hate speech classification. Performance metrics are all F1 scores. *Italics* indicates the worst performance. **Bold** face indicates best performance.

is a pejorative term used to describe Bernie Sanders supporters that have recently received criticism for crude and sexist attacks against rival Hillary Clinton. A similar example is shown in Example B of Table 7, where the term *Charlottesville* is used to refer to the white supremacist rally that took place in Charlottesville, Virginia rather than as a city’s name. Before post-processing, our model incorrectly predicts the tweet as a *threatening* post based on the words used in the post (*worry, fast*). By contrast, our model accurately detects the implicitly hateful message conveyed in the tweet after *Charlottesville* is being identified in its correct sense.

5.3 Challenges in Implicit Hate Speech Detection

Section 4.2 shows that Wikipedia knowledge does not add value and sometimes even hurts the implicit hate speech type classification. To further understand the role of Wikipedia knowledge in implicit hate type classification, we identify two challenges our model struggles with in implicit hate speech classification. One of the most common challenges for our model is shown in (3), where hatred is not apparent from the text but requires a pragmatic understanding of the text.

(3) **white privilege conference this week. #nationalsocialism white privilege, or white skin privilege, is the societal privilege that benefits white people over non-white people in some societies, particularly if they are otherwise un-**

der the same social, political, or economic circumstances. With roots in European colonialism and imperialism, and the Atlantic slave trade, white privilege has developed in circumstances that have broadly sought to protect white racial privileges, various national citizenships, and other rights or special benefits. (Gold label: Incitement, Prediction: White grievance)

(4) **atlanta hawks poor attendance is blamed on black crowds making southern whites uncomfortable. The Atlanta Hawks are a professional basketball team based in Atlanta, Georgia. The Hawks compete in the National Basketball Association (NBA) as a member team of the league’s Eastern Conference Southeast Division. The Hawks play their home games at Philips Arena.** (Gold: Incitement, Prediction: Stereotypical)

The post in (3) is an *incitement* post because it implicitly elevates white privilege ideology by promoting an upcoming event on white privilege. However, the tweet is predicted as a *white grievance* post because of the recurrence of the term *white privilege* in the post and Wikipedia description. Here we can see that world knowledge distracts the model. Additionally, our model struggles with this type of indirect hate because the implied hateful message is not apparent from the text. The model has to understand the pragmatic implicature (White people are privileged) that the

#	Gold	Prediction	Example
A	Explicit hate	Explicit hate	<i>the kkk. how old does my little boy need to be to join?</i> The Ku Klux Klan (KKK), or simply "the Klan", is the name of three distinct past and present movements in the United States that have advocated extremist reactionary currents such as white supremacy, white nationalism, and anti-immigration, historically expressed through terrorism aimed at groups or individuals whom they opposed.
B	White grievance	White grievance	<i>refuses to denounce dr.david duke! #makeamericagreatagain #votetrump</i> David Ernest Duke (born July 1, 1950) is an American white nationalist, anti-Semitic conspiracy theorist, politician, and former Grand Wizard of the Ku Klux Klan.
C	Incitement	Incitement	<i>#charlottesville a day that will go down in white history.</i> Charlottesville, a metonym for the Unite the Right rally, a white supremacist rally that took place in Charlottesville, Virginia, from August 11 to 12, 2017.

Table 6: Examples where Wikipedia description helps the model make the right decision. **Red** color indicates the features with highest coefficients for the model, darker colors indicate more polarity. The tweet is in *italics*. Wikipedia description is indicated in **bold**.

post suggests or implies in order to decipher the hidden hateful intent of the tweet.

(4) demonstrates another common challenge our model faces. The description of the basketball team does not help understand the underlying hateful-ness of the tweet, which implies that white people don't go to watch the basketball game because black players would make white people uncomfortable. That said, the Wikipedia description here does not hurt the model as well because the model still predicts the tweet as *stereotypical* without the description. The underlying problem here is that Wikipedia descriptions and tweet text seem to be unrelated to each other. This domain discrepancy between text and knowledge suggests that simple concatenation is not enough, but a more sophisticated structure that can capture the information flow between text and knowledge representation is needed for implicit hate speech type classification.

6 Conclusion

This paper has proposed the idea of integrating real world knowledge into the task of hate speech detection. Experimental results show that real world knowledge is helpful, especially in cases where entity triggers (e.g., rally, KKK) are present in the tweet. However, our analysis also shows that this knowledge fails to predict implicit hate speech types as Wikipedia knowledge does not add value and sometimes even hurts the classification, suggesting that a more sophisticated model that enables understanding beyond knowledge about the involved entities is required for implicit hate speech

type classification. To mitigate these challenges, works on model architecture that enable information flow between the representations of the tweet and Wikipedia knowledge is a reasonable next step. Additionally, exploring the possibility of combining different kinds of external knowledge, for example, combining commonsense knowledge (Chowdhury and Chaturvedi, 2021) in modeling implicit hate speech would also be a promising line of inquiry. To further understand the subtleties of hate speech, deciphering models for coded language or indirect language (e.g., metaphor, irony) in hate speech expression would be beneficial.

Ethics Statement

With the exponential growth of offensive language online, a myriad of machine learning models has been proposed. However, a major limitation of many existing hate speech detection models is that they focused on capturing explicit or overt hate speech, failing to detect implicit hateful expressions that are no less harmful than explicit ones. Our experiment in this study suggests that entity triggers (e.g., rally, KKK) are helpful in detecting hatefulness that is not apparent from the text. This could ideally help improve model accuracy in identifying implicit hate speech, preventing targeted communities from experiencing increased harm online. Furthermore, we show that knowledge about entities may help reduce false positives in explicit hate speech. This is important, as deep learning models nowadays still suffer from false positive predictions (Markov and Daelemans, 2021). To

#	Gold	Model predictions before and after post-processing	Example
A	Explicit hate	Non-hate	<i>bernie bros are entitled white males #maleprivilege #culturalmarxism #feminismiscancer #democrats</i> White is a racialized classification of people and a skin color specifier, generally used for people of European origin; although the definition can vary depending on context, nationality, and point of view. In the United States, this term has at times been expanded to encompass persons of Mexican, South Asian, West Asian, East Asian, and North African descent, persons who are often considered "non-White" in other contexts in the United States.
		Explicit hate	<i>bernie bros are entitled white males #maleprivilege #culturalmarxism #feminismiscancer #democrats</i> "Bernie Bro", or sometimes "Berniebro", is a pejorative term for a person (particularly one who is young, male and wealthy) who supports American politician and junior United States senator Bernie Sanders via Internet culture and social media by "mainstream" media outlets and later Hillary Clinton supporters. Recently, Bernie Bros have received criticism for crude and sexist attacks against rival Democratic nominee Hillary Clinton and her second presidential campaign in 2016, with her first attempt being in 2008.
B	Incitement	Threatening	<i>don't worry charlottesville was just the beginning. we're growing extremely fast.</i> Charlottesville, colloquially C'ville, Hoo-Ville, and formally the City of Charlottesville, is an independent city in the Commonwealth of Virginia. As of 2015, the population was 46,597.
		Incitement	<i>don't worry charlottesville was just the beginning. we're growing extremely fast.</i> Charlottesville, a metonym for the Unite the Right rally, a white supremacist rally that took place in Charlottesville, Virginia, from August 11 to 12, 2017.

Table 7: Ground truth labels along with predictions made by the model before and after post-processing. Red color indicates the features with the highest coefficients for the model, darker colors indicate more polarity. The tweet is in *italics*. Wikipedia description is indicated in **bold**.

this end, minimizing false positives is pivotal, as models that are not robust enough would be far from being applicable in the real world as a moderation tool, and using such a non-robust model would further lead to over-blocking or removal of harmless social media content that does not violate community guidelines.

References

- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Nozza Debora, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, Manuela Sanguinetti, et al. 2019. Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter. In *13th International Workshop on Semantic Evaluation*, pages 54–63. Association for Computational Linguistics.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. Comet: Commonsense transformers for automatic knowledge graph construction. *arXiv preprint arXiv:1906.05317*.
- Tommaso Caselli, Valerio Basile, Mitrović Jelena, Kar-toziya Inga, Granitzer Michael, et al. 2020. I feel offended, don't be abusive! implicit/explicit messages in offensive and abusive language. In *Language Resources and Evaluation Conference*, pages 6193–6202. The European Language Resources Association.
- Somnath Basu Roy Chowdhury and Snigdha Chaturvedi. 2021. Does commonsense help in detecting sarcasm? *arXiv preprint arXiv:2109.08588*.
- Thomas Davidson, Dana Warmesley, 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, pages 512–515.
- Mai ElSherief, Caleb Ziems, David Muchlinski, Vaishnavi Anupindi, Jordyn Seybolt, Munmun De Choudhury, and Diyi Yang. 2021. Latent hatred: A benchmark for understanding implicit hate speech. *arXiv preprint arXiv:2109.05322*.
- Antigoni Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large scale crowdsourcing and characterization of twitter abusive behavior. In *Twelfth International AAAI Conference on Web and Social Media*.
- Lei Gao and Ruihong Huang. 2017. Detecting online hate speech using context aware models. *arXiv preprint arXiv:1710.07395*.
- Svetlana Kiritchenko, Isar Nejadgholi, and Kathleen C Fraser. 2021. Confronting abusive language online: A survey from the ethical and human rights perspective. *Journal of Artificial Intelligence Research*, 71:431–478.
- Jiangnan Li, Hongliang Pan, Zheng Lin, Peng Fu, and Weiping Wang. 2021. Sarcasm detection with commonsense knowledge. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3192–3201.

- Ilija Markov and Walter Daelemans. 2021. [Improving cross-domain hate speech detection by reducing the false positive rate](#). In *Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda*, pages 17–22, Online. Association for Computational Linguistics.
- M Mozafari, R Farahbakhsh, and N Crespi. 2020. Hate speech detection and racial bias mitigation in social media based on bert model. *PLoS ONE*, 15(8):e0237861.
- Debaditya Pal, Kaustubh Chaudhari, and Harsh Sharma. 2022. Combating high variance in data-scarce implicit hate speech classification. *arXiv preprint arXiv:2208.13595*.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.
- Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A Smith, and Yejin Choi. 2019. Social bias frames: Reasoning about social and power implications of language. *arXiv preprint arXiv:1911.03891*.
- Johannes M van Hulst, Faegheh Hasibi, Koen Dercksen, Krisztian Balog, and Arjen P de Vries. 2020. Rel: An entity linker standing on the shoulders of giants. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2197–2200.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.
- Zeeraq Waseem, Thomas Davidson, Dana Warmusley, and Ingmar Weber. 2017. Understanding abuse: A typology of abusive language detection subtasks. *arXiv preprint arXiv:1705.09899*.
- Zeeraq 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.