
Interpretable Multi-Modal Hate Speech Detection

Prashanth Vijayaraghavan¹ Hugo Larochelle² Deb Roy¹

Abstract

With growing role of social media in shaping public opinions and beliefs across the world, there has been an increased attention to identify and counter the problem of hate speech on social media. Hate speech on online spaces has serious manifestations, including social polarization and hate crimes. While prior works have proposed automated techniques to detect hate speech online, these techniques primarily fail to look beyond the textual content. Moreover, few attempts have been made to focus on the aspects of interpretability of such models given the social and legal implications of incorrect predictions. In this work, we propose a deep neural multi-modal model that can: (a) detect hate speech by effectively capturing the semantics of the text along with socio-cultural context in which a particular hate expression is made, and (b) provide interpretable insights into decisions of our model. By performing a thorough evaluation of different modeling techniques, we demonstrate that our model is able to outperform the existing state-of-the-art approaches. Finally, we show the importance of social and cultural context features towards unearthing clusters associated with different categories of hate.

1. Introduction

Recently, Social media has become a breeding ground for abuse in the form of misuse of private user details (Yang & Srinivasan, 2014; Kontaxis et al., 2011), hate speech and cyberbullying (Djuric et al., 2015; Sood et al., 2012; Singh et al., 2017), and false or malicious information (Cheng et al., 2017; Aro, 2016; Weimann, 2010). Müller et al. (Müller & Schwarz, 2017; 2018) investigated the link between social media, specifically Facebook and Twitter and hate crime in countries like United States and Germany. The

¹MIT Media Lab, Cambridge, MA, USA ²Google Brain, Montreal, Canada. Correspondence to: Prashanth Vijayaraghavan <pralav@mit.edu>.

results of these studies suggest that social media can act as a propagation mechanism between online hate speech and real-life violent crimes. Therefore, the need for combating real-life hate crimes also involves improving the health of public discourse on such online platforms. The primary step in this process requires detecting and tracking hate online. This is also critical for protecting democracy from motivated malignancy. In this work, we, therefore, focus on hate speech – defined formally as any communication that disparages a person or a group on the basis of some characteristic such as race, color, ethnicity, gender, sexual orientation, nationality, religion, or other characteristic (Nockleby, 2000).

A recent study (Gröndahl et al., 2018; Hosseini et al., 2017) has exposed weaknesses in many machine learning detectors currently used to recognize and keep hate speech at bay. In addition to problems in dealing with language related challenges, prior work ignores an important aspect of hate speech – social and cultural dimension of hate speech. Hate speech covers a range of different type of expressions and it is critical to look at the context in which a particular expression is made: (a) What was the nature of expression? (b) Who made the expression? and (c) What was the socio-cultural context? Since language also evolves quickly, specifically among young users of social networks, it adds another layer of complexity (Nobata et al., 2016). Some communities tend to use benign words or phrases that have accepted hate speech meanings within their community and specific social context of usage. Therefore, we need to understand that hate speech may have strong socio-cultural implications (Raisi & Huang, 2016) that needs to be addressed.

Most of the current methods are trained primarily using the textual features containing limited examples of hate content and the results of their evaluations suggest that it is really difficult to detect hateful content without any additional contextual information (Kwok & Wang, 2013). In this work, we develop a hate speech classification model for tweets incorporating both social and cultural background of the user along with textual features. We further analyze the tweets to understand how the socio-political context helps us understand the categories of hate speech.

2. Dataset

One of the main challenges in building a robust hate speech detection model is the lack of a common benchmark dataset. Our dataset consists of tweets collected and aggregated from various different sources.

2.1. Publicly available datasets

There are different publicly available datasets usually obtained using keyword, hate lexicon or hashtag based filtering from Twitter stream over a period of time and then manually annotated. Table 1 gives details of all datasets used in our work. It is found that distinguishing hate speech from non-hate offensive language is a challenging task, as presence of offensive words is not a necessary condition for hate speech and offensive language does not always express hate.

2.2. Data collection

We further expand the existing dataset using data augmentation and distant supervision techniques. We collect additional data to add more capabilities to our modeling. We explain the steps involved in collecting these additional resources:

Tweets In order to expand our dataset, we perform an exploratory search on Twitter that consists of phrases containing: (a) swear words combined with positive adjectives (eg. “f**king awesome”, “damn good”, “bloody wonderful”) (b) swear words combined with races, religions, sexual orientations or derogatory references to them. (e.g. “f**king ragheads”, “sh**ty muslims”). We sample hundred random tweets from a list of tweets collected using (a) and (b) and manually annotated them. With high precision, the tweets obtained using (a) were non-hateful, while those tweets collected using (b) comprised of hate messages. This is backed by the procedure followed by Golbeck et al. (Golbeck et al., 2017; ElSherief et al., 2018) and subsequently manually annotated. We also get the hate code words as explained by Magu et al. (Magu et al., 2017) and obtain those tweets containing such code words along with swear words and identity based derogatory lexicons. We curate few random samples of this aggregated data using human annotators and majority of the annotation conformed with the classes they were assigned to. We perform data augmentation by introducing some noise in the form of common misspellings into the data.

Hate groups Collecting information about extremist groups and their followers can be useful in determining hate communities on Twitter. In this direction, we gather the data collected by Southern Poverty Law Center (SPLC)¹. SPLC

¹<https://www.splcenter.org/>

has been monitoring activities of various extremist groups in United States like anti-immigrant, white nationalists, the neo-Nazi movement, antigovernment militias and others. The SPLC website contains information about prominent members of such groups and their websites. In an effort to map these extremist groups on twitter, we perform a user account look up on Twitter by filtering those users who have the extremist group’s name as apart of their name or bio. We got $\sim 3k$ user accounts containing such information and filtering for inactive accounts. We construct a directed graph \mathcal{G} where each vertex is a user and edges represent their relationships (friends followers). We compute the page rank of this graph \mathcal{G} and obtain the top $\sim 10k$ accounts including the $\sim 3k$ seed user accounts. Figure 1 shows the the various categories of hate groups and summary of the the data collected. we construct a graph G^h using all the authors’ and “mentioned” users’ follower and friends list in our dataset and perform an intersection with the $\sim 10k$ hate community accounts.

Finally, we construct our dataset removing data with no user information. Table 1 gives the statistics of the aggregated dataset. We believe that all these resources can help us build an efficient hate speech classifier, considering language, social and cultural context as a whole.




Figure 1. Left: Percentage of seed user accounts belonging to a particular hate category. Right: Sample user accounts related to each hate category.

3. Model

In this section, we describe our deep learning model that accounts for the multimodal nature of our data. The data

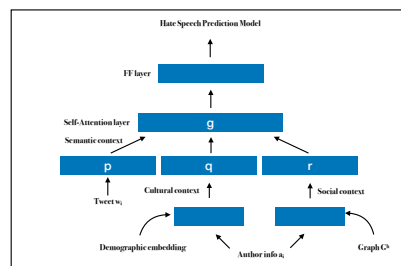


Figure 2. Illustration of the model

Datasets	Details
(Founta et al., 2018)	None: 53.8%; Hate: 4.96%; Abusive: 27.15%; Spam: 14%; Tweets: $\sim 100k$;
(Davidson et al., 2017)	None: 16.8%; Hate: 5.8%; Offensive: 77.4%; Tweets: $\sim 25k$;
(Park & Fung, 2017)	None: 68%; Sexism: 20%; Racism: 11%; Tweets: $\sim 18k$;
(Golbeck et al., 2017)	None: 74%; Harassment: 26%; Tweets: $\sim 21k$;
Our dataset	None: 58.1%; Hate: 16.6%; Abusive: 25.3% Tweets: $\sim 258k$;

Table 1. Summary of different datasets.

includes tweets, author attributes and social structure represented by author’s followers and friends. Let us denote our hate dataset as $D^{(H)} = (w_1, a_1), (w_2, a_2), \dots, (w_n, a_n)$, where each tuple in this set consists of the tweet text w_i , author information a_i used to derive social and cultural context associated with the tweet. The dataset $D^{(H)}$ is divided into $D_{train}^{(H)}$, $D_{val}^{(H)}$ and $D_{test}^{(H)}$ for training, validation and testing purposes. Defining the input as $x_i = (w_i, a_i)$, we denote our model as:

$$f_{\theta}(x_i) \approx g_{\theta}(P(w_i), Q(a_i), R(a_i)) \quad (1)$$

where P, Q, R are neural architectures extracting semantic features from tweet text (P) and socio-cultural features (R, Q) from author information; g is the function that determines the fusion strategy. Figure 2 is an illustration of our model.

3.1. Extracting semantic features (p)

There are a number of recent works (Kim, 2014; Kiros et al., 2015; Conneau et al., 2017) that have focused on encoding the text into useful continuous feature representations. Our choice of the text encoder is motivated by the need to encode character information along with the word information to handle noisy nature of tweets. Our encoder is a slight variant of Chen et al. (Chen et al., 2018). This approach providing multiple level of granularity can be useful in order to handle rare or noisy words in the text. Given character embeddings $E^{(c)} = [e_1^{(c)}, e_2^{(c)}, \dots, e_{n'}^{(c)}]$ and word embeddings $E^{(w)} = [e_1^{(w)}, e_2^{(w)}, \dots, e_n^{(w)}]$ of the input, starting (p_t) and ending (q_t) character positions at time step t , we define inside character embeddings as: $E_I^{(c)} = [e_{p_t}^{(c)}, \dots, e_{q_t}^{(c)}]$ and outside embeddings as: $E_O^{(c)} = [e_1^{(c)}, \dots, e_{p_t-1}^{(c)}; e_{q_t+1}^{(c)}, \dots, e_{n'}^{(c)}]$. First, we obtain the character-enhanced word representation h_t by combining the word information from $E^{(w)}$ with the character context vectors. Character context vectors are obtained by attending over inside and outside character embeddings. Next, we compute a summary vector S over the hidden states h_t using a self-attention mechanism expressed as $Attn(H)$. The summary vector S is the final textual feature representation. The attention mechanism helps us to recognize the most relevant text units that contribute towards hate speech prediction.

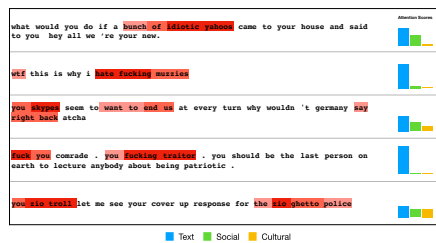


Figure 3. Table shows examples with attention scores of textual, social and cultural context features. Text containing hate code words have higher attention scores (highlighted) for social and cultural context features.

3.2. Extracting cultural context features (q)

The demographic features act as a proxy for deriving the cultural context information required for our classification task. A work by Vijayaraghavan et al. (Vijayaraghavan et al., 2017) developed a multimodal demographic classifier for Twitter users. By feeding author information a_i to this model, we extract the penultimate layer demographic representations for all users in our dataset.

3.3. Extracting social context features (r)

Social context is obtained using the graph G^h explained in Section 2. We use an approach delineated by Vijayaraghavan et al. (Vijayaraghavan et al., 2017), we build a binary vector for each tweet’s author a_i based on the $\sim 10k$ hate community accounts. Each index of the vector represents one of those hate group accounts the value (1, 0) represents if the author is following the account or not. We use feed-forward layers to map this binary vector into a lower dimension vector.

3.3.1. COMBINING TEXTUAL AND SOCIO-CULTURAL FEATURES (G)

We apply a late fusion strategy in our work. The multimodal features are fed through a self-attention mechanism to understand the relevance of each of these different features towards the hate speech classification task. We perform a weighted sum of these feature vectors to obtain a final vector representation R^h . We use categorical cross-entropy as our

Model	F1 (Hate)	F1 (Overall)
Traditional Models: Text+Social+Demographic		
LR	0.53	0.72
SVM	0.563	0.729
DL Models: Text Only		
CNN-Char	0.735	0.866
BiGRU+Char+Attn	0.744	0.864
CNN-Word	0.658	0.788
BiGRU+Attn	0.683	0.801
BiLSTM-2DCNN	0.661	0.795
DL Models: Text+SC		
CNN-Char+FF	0.760	0.879
BiGRU+Char+Attn+FF	0.784	0.90

Table 2. Results of our experiments.

objective to train our model.

4. Evaluation

We experimented with different models. Table 2 shows the results of our evaluations using traditional and other baseline deep learning models. It is clear that deep learning models outperform the traditional models. Table 2 also shows that social and cultural context improves the performance of our model significantly.

4.0.1. IDENTIFYING CATEGORIES OF HATE SPEECH

We verify if the embeddings learned from the model can help identify different hate categories. To achieve that, we compute the final representation R^h before the final feed-forward layer for every tweet classified as hate speech by our model. This representation R^h is our hate embedding. We use a simple agglomerative clustering method on our hate embeddings with a parameter $k = 5$ for the number of clusters. We manually label these clusters using the top words ranked based on attention scores. Table 3 shows the top 5 words associated with the clusters and manual hate category assigned to them.

Top 5 words	Hate Category
jihadi, muzzie, terrorist, #stopislam, #banmuslim	Anti-Islam
n**ga, n**ger, #whitepower, ghetto #14words	Anti-Black
#buildthewall, illegals, #noamnesty, #illegalaliens, #anchorbabies	Anti-Immigrant
f**k, c**t, hate, b**ch, a****le	General Hate
#antisemitism, #antisemites, nazi, satan, neonazi	Anti-semitic

Table 3. Top 5 words associated with different clusters and manually assigned hate category.

We evaluate the relevance of social and cultural context features by comparing the clusters produced by embeddings obtained from two models: Text-Only Model (BiGRU+Char+Attn) and Text+SC Model (BiGRU+Char+Attn+FF). We measure the ability of these

models to recover the hate categories with our embeddings through a purity score. We measure the amount of overlap between two clusterings as: $Purity = \frac{1}{N} \sum_i \max_j |g_i \cap c_j|$, where N is the total number of input data points, g_i is the i -th ground truth hate category, and c_j is the j -th predicted cluster.

Since we do not have the ground truth hate category for our tweets, we sampled 100 random tweets from every cluster obtained from both the models and obtained human judgments. Workers were shown the original tweet and asked the hate category they fall under. They were also allowed to choose 'None of the above' as an option. Once we obtain the ground truths using these human annotations, we computed the purity scores. The results are shown in Table 4. The results indicate that the model that uses social and cultural context is able to produce better clusters having more overlap with ground truth hate categories compared to the text only model. The performance can be attributed to the Hate graph G^h containing inherent information of various categories of hate groups.

Model	Purity Score
Text Only	0.52
Text+SC	0.76

Table 4. Cluster purity scores

4.1. Interpretability

In order to be able to draw insights on the decisions of the model, we qualitatively look at the data where the social context information is used for prediction. In Section 3, we explained the need for an attention layer that fuses both the semantic text features and social and demographic features. In order to assess the importance of social and cultural context for hate speech classification, we identify examples in test set that has a high attention score for social context vector. We visualize by highlighting words in the text and constructing bar graphs that indicate the relevance of each of the features: textual, social, cultural). We compared the results from attention scores with perturbations-based approach for interpretability. Since they were strikingly similar, we show the results based on attention scores only. Figure 3 shows few examples where that is the case. Interestingly, these are tweets containing code words like "skypes", "yahoos", "zio", "zog", etc. attacking particular group of people. These kinds of code words are widespread among the hate groups (explained in Section 2). Though our dataset contains only few instances of such words in the presence of swear words, the model is still able to understand these code words and tag them as hateful content which otherwise could just pass on as being abusive but not hateful (though it can be debatable).

5. Conclusion

In this work, we developed a comprehensive model to automatically detect hateful content. Motivated by the need to incorporate socio-cultural information for hate speech detection as a part of the hate speech detection task, we adopt different feature extraction strategies for different modalities of data: text, demographic information and social graph. We build a multi-modal classification method with late-fusion of different modalities of data. We derive important insights about our model and its ability to understand hate speech code words and cluster into different categories of hate speech.

References

- Aro, J. The cyberspace war: propaganda and trolling as warfare tools. *European View*, 15(1):121–132, 2016.
- Chen, H., Huang, S., Chiang, D., Dai, X., and Chen, J. Combining character and word information in neural machine translation using a multi-level attention. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pp. 1284–1293, 2018.
- Cheng, J., Bernstein, M., Danescu-Niculescu-Mizil, C., and Leskovec, J. Anyone can become a troll: Causes of trolling behavior in online discussions. In *CSCW: proceedings of the Conference on Computer-Supported Cooperative Work. Conference on Computer-Supported Cooperative Work*, volume 2017, pp. 1217. NIH Public Access, 2017.
- Conneau, A., Kiela, D., Schwenk, H., Barrault, L., and Bordes, A. Supervised learning of universal sentence representations from natural language inference data. *arXiv preprint arXiv:1705.02364*, 2017.
- Davidson, T., Warmesley, D., Macy, M., and Weber, I. Automated hate speech detection and the problem of offensive language. *arXiv preprint arXiv:1703.04009*, 2017.
- Djuric, N., Zhou, J., Morris, R., Grbovic, M., Radosavljevic, V., and Bhamidipati, N. Hate speech detection with comment embeddings. In *Proceedings of the 24th international conference on world wide web*, pp. 29–30. ACM, 2015.
- ElSherief, M., Kulkarni, V., Nguyen, D., Wang, W. Y., and Belding, E. Hate lingo: A target-based linguistic analysis of hate speech in social media. In *Proceedings of the 12th International AAAI Conference on Web and Social Media, ICWSM '18*, 2018.
- Founta, A.-M., Djouvas, C., Chatzakou, D., Leontiadis, I., Blackburn, J., Stringhini, G., Vakali, A., Sirivianos, M., and Kourtellis, N. Large scale crowdsourcing and characterization of twitter abusive behavior. In *11th International Conference on Web and Social Media, ICWSM 2018*. AAAI Press, 2018.
- Golbeck, J., Ashktorab, Z., Banjo, R. O., Berlinger, A., Bhagwan, S., Buntain, C., Chekalos, P., Geller, A. A., Gergory, Q., Gnanasekaran, R. K., et al. A large labeled corpus for online harassment research. In *Proceedings of the 2017 ACM on Web Science Conference*, pp. 229–233. ACM, 2017.
- Gröndahl, T., Pajola, L., Juuti, M., Conti, M., and Asokan, N. All you need is” love”: Evading hate-speech detection. *arXiv preprint arXiv:1808.09115*, 2018.
- Hosseini, H., Kannan, S., Zhang, B., and Poovendran, R. Deceiving google’s perspective api built for detecting toxic comments. *arXiv preprint arXiv:1702.08138*, 2017.
- Kim, Y. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.
- Kiros, R., Zhu, Y., Salakhutdinov, R. R., Zemel, R., Urtasun, R., Torralba, A., and Fidler, S. Skip-thought vectors. In *Advances in neural information processing systems*, pp. 3294–3302, 2015.
- Kontaxis, G., Polakis, I., Ioannidis, S., and Markatos, E. P. Detecting social network profile cloning. In *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2011 IEEE International Conference on*, pp. 295–300. IEEE, 2011.
- Kwok, I. and Wang, Y. Locate the hate: Detecting tweets against blacks. In *AAAI*, 2013.
- Magu, R., Joshi, K., and Luo, J. Detecting the hate code on social media. *arXiv preprint arXiv:1703.05443*, 2017.
- Müller, K. and Schwarz, C. Fanning the flames of hate: Social media and hate crime. 2017.
- Müller, K. and Schwarz, C. Making america hate again? twitter and hate crime under trump. 2018.
- Nobata, C., Tetreault, J., Thomas, A., Mehdad, Y., and Chang, Y. Abusive language detection in online user content. In *Proceedings of the 25th international conference on world wide web*, pp. 145–153. International World Wide Web Conferences Steering Committee, 2016.
- Nockleby, J. T. Hate speech. *Encyclopedia of the American constitution*, 3:1277–79, 2000.
- Park, J. H. and Fung, P. One-step and two-step classification for abusive language detection on twitter. *arXiv preprint arXiv:1706.01206*, 2017.
- Raisi, E. and Huang, B. Cyberbullying identification using participant-vocabulary consistency. *arXiv preprint arXiv:1606.08084*, 2016.
- Singh, V. K., Radford, M. L., Huang, Q., and Furrer, S. ” they basically like destroyed the school one day”: On newer app features and cyberbullying in schools. In *CSCW*, pp. 1210–1216, 2017.
- Sood, S., Antin, J., and Churchill, E. Profanity use in online communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1481–1490. ACM, 2012.
- Vijayaraghavan, P., Vosoughi, S., and Roy, D. Twitter demographic classification using deep multi-modal multi-task learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pp. 478–483, 2017.
- Weimann, G. Terror on facebook, twitter, and youtube. *The Brown Journal of World Affairs*, 16(2):45–54, 2010.
- Yang, C. and Srinivasan, P. Translating surveys to surveillance on social media: methodological challenges & solutions. In *Proceedings of the 2014 ACM conference on Web science*, pp. 4–12. ACM, 2014.