

Detecting Online Hate Speech Using Context Aware Models

Lei Gao

Texas A&M University
sjtuprog@tamu.edu

Ruihong Huang

Texas A&M University
huangrh@cse.tamu.edu

Abstract

In the wake of a polarizing election, the cyber world is laden with hate speech. Context accompanying a hate speech text is useful for identifying hate speech, which however has been largely overlooked in existing datasets and hate speech detection models. In this paper, we provide an annotated corpus of hate speech with context information well kept. Then we propose two types of hate speech detection models that incorporate context information, a logistic regression model with context features and a neural network model with learning components for context. Our evaluation shows that both models outperform a strong baseline by around 3% to 4% in F1 score and combining these two models further improve the performance by another 7% in F1 score.

1 Introduction

Following a turbulent election season, 2016’s cyber world is awash with hate speech. Automatic detection of hate speech has become an urgent need since human supervision is unable to deal with large quantities of emerging texts.

Context information, by our definition, is the text, symbols or any other kind of information related to the original text. While intuitively, context accompanying hate speech is useful for detecting hate speech, context information of hate speech has been overlooked in existing datasets and automatic detection models.

Online hate speech tends to be subtle and creative, which makes context especially important for automatic hate speech detection. For instance,

(1) *barryswallows: Merkel would never say NO*

This comment is posted for the news titled by “German lawmakers approve ‘no means no’ rape law after Cologne assaults”. With context, it becomes clear that this comment is a vicious insult towards female politician. However, almost all the publicly available hate speech annotated datasets do not contain context information. [Waseem and Hovy \(2016\)](#); [Waseem \(2016\)](#); [Wulczyn et al. \(2016\)](#); [Ross et al. \(2017\)](#).

We have created a new dataset consisting of 1528 Fox News user comments, which were taken from 10 complete discussion threads for 10 widely read Fox News articles. It is different from previous datasets from the following two perspectives. First, it preserves rich context information for each comment, including its user screen name, all comments in the same thread and the news article the comment is written for. Second, there is no biased data selection and all comments in each news comment thread were annotated.

In this paper, we explored two types of models, feature based logistic regression models and neural network models, in order to incorporate context information in automatic hate speech detection. First, logistic regression models have been used in several prior hate speech detection studies [Chen et al. \(2012\)](#); [Burnap and Williams \(2014\)](#); [Van Hee et al. \(2015\)](#); [Hosseinmardi et al. \(2015\)](#); [Burnap and Williams \(2015\)](#); [Waseem and Hovy \(2016\)](#); [Wulczyn et al. \(2016\)](#); [Nobata et al. \(2016\)](#) and various features have been tried including character-level and word-level n-gram features, syntactic features, linguistic features, and comment embedding features. However, all the features were derived from the to-be-classified text itself. In contrast, we experiment with logistic regression models using features extracted from context text as well. Second, neural network models [Zhang et al. \(2015\)](#); [Tang et al. \(2015\)](#); [Yang et al. \(2016\)](#) have the potential to capture compo-

sitional meanings of text, but they have not been well explored for online hate speech detection until recently [Pavlopoulos et al. \(2017\)](#). We experiment with neural net models containing separate learning components that model compositional meanings of context information. Furthermore, recognizing unique strengths of each type of models, we build ensemble models of the two types of models. Evaluation shows that context-aware logistic regression models and neural net models outperform their counterparts that are blind with context information. Especially, the final ensemble models outperform a strong baseline system by around 10% in F1-score.

2 Related Works

Recently, a few datasets with human labeled hate speech have been created, however, most of existing datasets do not contain context information. Due to the sparsity of hate speech in everyday posts, researchers tend to sample candidates from bootstrapping instead of random sampling, in order to increase the chance of seeing hate speech. Therefore, the collected data instances are likely to be from distinct contexts.

For instance, in the Primary Data Set described in [Djuric et al. \(2015\)](#) and later used by [Nobata et al. \(2016\)](#), 10% of the dataset is randomly selected while the remaining consists of comments tagged by users and editors. [Kwok and Wang \(2013\)](#) built a balanced data set of 24.5k tweets by selecting from Twitter accounts that claimed to be racist or were deemed racist using their followed news sources. [Burnap and Williams \(2014\)](#) collected hateful tweets related to the murder of Drummer Lee Rigby in 2013. [Waseem and Hovy \(2016\)](#) provided a corpus of 16k annotated tweets in which 3.3k are labeled as sexist and 1.9k are labeled as racist. They created this corpus by bootstrapping from certain key words ,specific hashtags and certain prolific users. [Warner and Hirschberg \(2012\)](#) created a dataset of 9000 human labeled paragraphs that were collected using regular expression matching in order to find hate speech targeting Judaism and Israel. [Hosseinmardi et al. \(2015\)](#) extracted data instances from instagram that were associated with certain user accounts. [Wulczyn et al. \(2016\)](#) presented a very large corpus containing over 115k wikipedia comments that include around 37k randomly sampled comments and the remaining 78k comments were

selected from wikipedia blocked comments.

Most of existing hate speech detection models are feature based and use features derived from the target text itself. [Burnap and Williams \(2014\)](#) experimented with different classification methods including Bayesian Logistic Regression, Random Forest Decision Trees and SVMs, using features such as n-grams, reduced n-grams, dependency paths, and hateful terms. [Waseem and Hovy \(2016\)](#) proposed a logistic regression model using character n-gram features. [Djuric et al. \(2015\)](#) used the paragraph2vec for joint modeling of comments and words, then the generated embeddings were used as feature in a logistic regression model. [Nobata et al. \(2016\)](#) experimented with various syntactic, linguistic and distributional semantic features including word length, sentence length, part of speech tags, and embedding features, in order to improve performance of logistic regression classifiers. Recently, [Schmidt and Wiegand \(2017\)](#) surveyed current approaches for hate speech detection, which interestingly also called to attention on modeling context information for resolving difficult hate speech instances.

3 The Fox News User Comments corpus

3.1 Corpus Overview

The Fox News User Comments corpus consists of 1528 annotated comments (435 labeled as hateful) that were posted by 678 different users in 10 complete news discussion threads in the Fox News website. The 10 threads were manually selected and represent popular discussion threads during August 2016. All of the comments included in these 10 threads were annotated. The number of comments in each of the 10 threads is roughly equal. Rich context information was kept for each comment, including its user screen name, the comments and their nested structure and the original news article. The data corpus along with annotation guidelines is posted on github¹.

3.2 Annotation Guidelines

Our annotation guidelines are similar to the guidelines used by [Nobata et al. \(2016\)](#). We define hateful speech to be the language which explicitly or implicitly threatens or demeans a person or a group based upon a facet of their identity such as gender, ethnicity, or sexual orientation. The la-

¹<https://github.com/sjtuprog/fox-news-comments>

belonging of hateful speech in our corpus is binary. A comment will be labeled as hateful or non-hateful.

3.3 Annotation Procedure

We identified two native English speakers for annotating online user comments. The two annotators first discussed and practiced before they started annotation. They achieved a surprisingly high Kappa score Cohen (1960) of 0.98 on 648 comments from 4 threads. We think that thorough discussions in the training stage is the key for achieving this high inter-agreement. For those comments which annotators disagreed on, we label them as hateful as long as one annotator labeled them as hateful. Then one annotator continued to annotate the remaining 880 comments from the remaining six discussion threads.

3.4 Characteristics in Fox News User Comments corpus

Hateful comments in the Fox News User Comments Corpus is often subtle, creative and implicit. Therefore, context information is necessary in order to accurately identify such hate speech.

3.4.1 Context Dependent Comments

The hatefulness of many comments depended on understanding their contexts. For instance,

(3) *mastersundholm: Just remember no trabjo no cervesa*

This comment is posted for the news "States moving to restore work requirements for food stamp recipients". This comment implies that Latino immigrants abuse the usage of food stamp policy, which is clearly a stereotyping.

3.4.2 Implicit and creative language

Many hateful comments use implicit and subtle language, which contain no clear hate indicating word or phrase. In order to recognize such hard cases, we hypothesize that neural net models are more suitable by capturing overall composite meanings of a comment. For instance, the following comment is a typical implicit stereotyping against women.

(4) *MarineAssassin: Hey Brianne - get in the kitchen and make me a samich. Chop Chop*

3.4.3 Long Comments with Regional Focus of hatefulness

11% of our annotated comments have more than 50 words each. In such long comments, the hate-

ful indicators usually appear in a small region of a comment while the majority of the comment is neutral. For example,

(5) *TMmckay: I thought ...115 words... Too many blacks winning, must be racist and needs affirmative action to make whites equally win!*

3.4.4 Disrespectful screen names

Certain user screen names indicate hatefulness, which imply that comments posted by these users are likely to contain hate speech. In the following example, commie is a slur for communists.

(6) *nocommie11: Blah blah blah. Israel is the only civilized nation in the region to keep the unwashed masses at bay.*

4 Context-aware Online Hate Speech Detection Models

4.1 Logistic Regression Models

In logistic regression models, we extract four types of features, word-level and character-level n-gram features as well as two types of lexicon derived features. We extract these four types of features from the target comment first. Then we extract these features from two sources of context texts, specifically the title of the news article that the comment was posted for and the screen name of the user who posted the comment.

For logistic regression model implementation, we use l2 loss. We adopt the balanced class weight as described in Scikit learn². Logistic regression model with character-level n-gram features is presented as a strong baseline for comparison since it was shown very effective. (Waseem and Hovy, 2016; Nobata et al., 2016)

4.1.1 Word-level and Character-level N-gram Features

For character level n-grams, we extract character level bigrams, tri-grams and four-grams. For word level n-grams, we extract unigrams and bigrams.

4.1.2 LIWC Feature

Linguistic Inquiry and Word Count, also called LIWC, has been proven useful for text analysis and classification Pennebaker et al. (2001). In the LIWC dictionary, each word is labeled with several semantic labels. In our experiment, we use

²http://scikit-learn.org/stable/modules/generated/sklearn.linear_model.LogisticRegression.html

the LIWC 2015 dictionary which contain 125 semantic categories. Each word is converted into a 125 dimension LIWC vector, one dimension per semantic category. The LIWC feature vector for a comment or its context is a 125 dimension vector as well, which is the sum of all its words' LIWC vectors.

4.1.3 NRC Emotion Lexicon Feature

NRC emotion lexicon contains a list of English words that were labeled with eight basic emotions (anger, fear, anticipation, trust, surprise, sadness, joy, and disgust) and sentiment polarities (negative and positive)([Mohammad and Turney, 2013](#)). We use NRC emotion lexicon to capture emotion clues in text. Each word is converted into a 10 dimension emotion vector, corresponding to eight emotion types and two polarity labels. The emotion vector for a comment or its context is a 10 dimension vector as well, which is the sum of all its words' emotion vectors.

4.2 Neural Network Models

Our neural network model mainly consists of three parallel LSTM [Hochreiter and Schmidhuber \(1997\)](#) layers. It has three different inputs, including the target comment, its news title and its username. Comment and news title are encoded into a sequence of word embeddings. We use pre-trained word embeddings in word2vec³. Username is encoded into a sequence of characters. We use one-hot encoding of characters.

Comment is sent into a bi-directional LSTM with attention mechanism. ([Bahdanau et al., 2014](#)). News title and username are sent into a bi-directional LSTM. Note that we did not apply attention mechanism to the neural network models for username and news title because both types of context are relatively short and attention mechanism tends to be useful when text input is long. The three LSTM output layers are concatenated, then connected to a sigmoid layer, which outputs predictions.

The number of hidden units in each LSTM used in our model is set to be 100. The recurrent dropout rate of LSTMs is set to 0.2. In addition, we use binary cross entropy as the loss function and a batch size of 128. The neural network models are trained for 30 epochs.

³<https://code.google.com/archive/p/word2vec/>

4.3 Ensemble Models

To study the difference of logistic regression model and neural network model and potentially get performance improvement, we will build and evaluate ensemble models.

5 Evaluation

We evaluate our model by 10 fold cross validation using our newly created Fox News User Comments Corpus. Both types of models use the exact same 10 folds of training data and test data. We report experimental results using multiple metrics, including accuracy, precision/recall/F1-score, and accuracy area under curve (AUC).

5.1 Experimental Results

5.1.1 Logistic Regression Models

Table 1 shows the performance of logistic regression models. The first section of table 1 shows the performance of logistic regression models using features extracted from a target comment only. The result shows that the logistic regression model was improved in every metric after adding both word-level n-gram features and lexicon derived features. However, the improvements are moderate.

The second section shows the performance of logistic regression models using the four types of features extracted from both a target comment and its contexts. The result shows that the logistic regression model using features extracted from a comment and both types of context achieved the best performance and obtained improvements of 2.8% and 2.5% in AUC score and F1-score respectively.

5.1.2 Neural Network Models

Table 2 shows the performance of neural network models. The first section of table 2 shows the performance of several neural network models that use comments as the only input. The model names are self-explanatory. We can see that the attention mechanism coupled with the bi-directional LSTM neural net greatly improved the online hate speech detection, by 5.7% in AUC score.

The second section of table 2 shows performance of the best neural net model (bi-directional LSTM with attention) after adding additional learning components that take context as input. The results show that adding username and news title can both improve model performance. Using

| Features | Input Contents | Accuracy | Precision | Recall | F1 | AUC |
|-----------------|-----------------------------------|--------------|--------------|--------------|--------------|--------------|
| char (baseline) | comment | 0.738 | 0.549 | 0.469 | 0.504 | 0.733 |
| +word | comment | 0.735 | 0.548 | 0.443 | 0.488 | 0.736 |
| +LIWC+NRC | comment | 0.732 | 0.533 | 0.465 | 0.495 | 0.740 |
| +word+LIWC+NRC | comment | 0.747 | 0.568 | 0.476 | 0.517 | 0.750 |
| | + username | 0.747 | 0.576 | 0.474 | 0.518 | 0.765 |
| | + title | 0.745 | 0.558 | 0.496 | 0.523 | 0.761 |
| | + title+ username (Best) | 0.750 | 0.572 | 0.516 | 0.542 | 0.778 |

Table 1: Performance of Logistic Regression Models

| Model | Input Contents | Accuracy | Precision | Recall | F1 | AUC |
|------------------------|-------------------------|--------------|--------------|--------------|--------------|--------------|
| LSTM | comment | 0.726 | 0.524 | 0.398 | 0.450 | 0.678 |
| bi-LSTM | comment | 0.720 | 0.513 | 0.440 | 0.473 | 0.682 |
| bi-LSTM with attention | comment | 0.750 | 0.591 | 0.437 | 0.499 | 0.735 |
| | + username | 0.742 | 0.566 | 0.437 | 0.489 | 0.748 |
| | + title (best) | 0.766 | 0.614 | 0.499 | 0.548 | 0.760 |
| | + title + username | 0.755 | 0.589 | 0.496 | 0.532 | 0.766 |

Table 2: Performance of Neural Network Models

news title gives the best F1 score while using both news title and username gives the best AUC score.

5.1.3 Ensemble Models

Table 3 shows performance of ensemble models by combining prediction results of the best context-aware logistic regression model and the best context-aware neural network model. We used two strategies in combining prediction results of two types of models. Specifically, the Max Score Ensemble model made the final decisions based on the maximum of two scores assigned by the two separate models; instead, the Average Score Ensemble model used the average score to make final decisions.

We can see that both ensemble models further improved hate speech detection performance compared with using one model only and achieved the best classification performance. Compared with the logistic regression baseline, the Max Score Ensemble model improved the recall by more than 20% with a comparable precision and improved the F1 score by around 10%, in addition, the Average Score Ensemble model improved the AUC score by around 7%.

6 Analysis

6.1 Logistic Regression Models

As shown in table 1, given comment as the only input content, the combination of character n-grams, word n-grams, LIWC feature and NRC feature achieves the best performance. It shows that in

addition to character level features, adding more features can improve hate speech detection performance. However, the improvement is limited. Compared with baseline model, the F1 score only improves 1.3%.

In contrast, when context information was taken into account, the performance greatly improved. Specifically, after incorporating features extracted from the news title and username, the model performance was improved by around 4% in both F1 score and AUC score. This shows that using additional context based features in logistic regression models is useful for hate speech detection.

6.2 Neural Network Models

As shown in table 2, given comment as the only input content, the bi-directional LSTM model with attention mechanism achieves the best performance. Note that the attention mechanism significantly improves the hate speech detection performance of the bi-directional LSTM model. We hypothesize that this is because hate indicator phrases are often concentrated in a small region of a comment, which is especially the case for long comments.

6.3 Ensemble Models

As shown in table 3, both ensemble models significantly improved hate speech detection performance. Figure 1 shows the system prediction results of comments that were labeled as hateful in the dataset. It can be seen that the two models per-

| Model | Accuracy | Precision | Recall | F1 | AUC |
|--------------------------------|--------------|--------------|--------------|--------------|--------------|
| Char (Baseline) | 0.738 | 0.549 | 0.469 | 0.504 | 0.733 |
| Best Neural Network Model | 0.766 | 0.614 | 0.499 | 0.548 | 0.760 |
| Best Logistic Regression Model | 0.750 | 0.572 | 0.516 | 0.542 | 0.778 |
| Max Score Ensemble | 0.740 | 0.539 | 0.678 | 0.600 | 0.794 |
| Average Score Ensemble | 0.779 | 0.650 | 0.496 | 0.560 | 0.804 |

Table 3: Performance of Ensemble Models

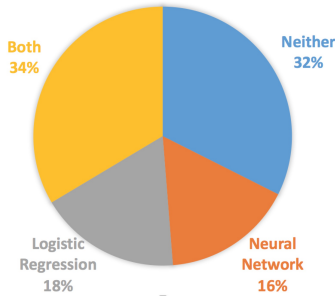


Figure 1: System Prediction Results of Comments that were Annotated as Hateful

form differently. We further examined predicted comments and find that both types of models have unique strengths in identifying certain types of hateful comments.

6.3.1 Strengths of Logistic Regression Models

The feature based logistic regression models are capable of making good use of character-level n-gram features, which are powerful in identifying hateful comments that contains OOV words, capitalized words or misspelled words. We provide two examples from the hateful comments that were only labeled by the logistic regression model:

(7)*kmawhmf:FBLM.*

Here FBLM means fuck Black Lives Matter. This hateful comment contains only character information which can exactly be made use of by our logistic regression model.

(8)*SFgunrnmn: what a efen loon, but most fem-anazis are.*

This comment deliberately misspelled feminazi for femanzis, which is a derogatory term for feminists. It shows that logistic regression model is capable in dealing with misspelling.

6.3.2 Strengths of Neural Network Models

The LSTM with attention mechanism are suitable for identifying specific small regions indicating hatefulness in long comments. In addition, the

neural net models are powerful in capturing implicit hateful language as well. The following are two hateful comment examples that were only identified by the neural net model:

(9)*freedomscout: @LarJass Many religions are poisonous to logic and truth, that much is true...and human beings still remain fallen human beings even they are Redeemed by the Sacrifice of Jesus Christ. So there's that. But the fallacies of thinking cannot be limited or attributed to religion but to error inherent in human motivation, the motivation to utter self-centeredness as fallen sinful human beings. Nearly all of the world's many religions are expressions of that utter sinful nature...Christianity and Judaism being the sole exceptions.*

This comment is expressing the stereotyping against religions which are not Christian or Judaism. The hatefulness is concentrated within the two bolded segments.

(10)*mamahattheridge: blacks Love being victims.*

In this comment, the four words themselves are not hateful at all. But when combined together, it is clearly hateful against black people.

7 Conclusion

We demonstrated the importance of utilizing context information for online hate speech detection. We first presented a corpus of hateful speech consisting of full threads of online discussion posts. In addition, we presented two types of models, feature based logistic regression models and neural network models, in order to incorporate context information for improving hate speech detection performance. Furthermore, we show that ensemble models leveraging strengths of both types of models achieve the best performance for automatic online hate speech detection.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- 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.
- Peter Burnap and Matthew Leighton Williams. 2014. Hate speech, machine classification and statistical modelling of information flows on twitter: Interpretation and communication for policy decision making. In *Proceedings of the Internet, Politics, and Policy conference*.
- Ying Chen, Yilu Zhou, Sencun Zhu, and Heng Xu. 2012. Detecting offensive language in social media to protect adolescent online safety. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Conference on Social Computing (SocialCom)*. IEEE, pages 71–80.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement* 20(1):37–46.
- Nemanja Djuric, Jing Zhou, Robin Morris, Mihajlo Grbovic, Vladan Radosavljevic, and Narayan Bhamidipati. 2015. Hate speech detection with comment embeddings. In *Proceedings of the 24th International Conference on World Wide Web*. ACM, pages 29–30.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Homa Hosseinmardi, Sabrina Arredondo Mattson, Rahat Ibn Rafiq, Richard Han, Qin Lv, and Shivakant Mishra. 2015. Detection of cyberbullying incidents on the instagram social network. *arXiv preprint arXiv:1503.03909*.
- Irene Kwok and Yuzhou Wang. 2013. Locate the hate: Detecting tweets against blacks. In *AAAI*.
- Saif M Mohammad and Peter D Turney. 2013. Nrc emotion lexicon. Technical report, NRC Technical Report.
- Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive language detection in online user content. In *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, pages 145–153.
- John Pavlopoulos, Prodromos Malakasiotis, and Ion Androutsopoulos. 2017. Deep learning for user comment moderation. *arXiv preprint arXiv:1705.09993*.
- James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic inquiry and word count: Liwc 2001. *Mahway: Lawrence Erlbaum Associates* 71(2001):2001.
- Björn Ross, Michael Rist, Guillermo Carbonell, Benjamin Cabrera, Nils Kurovsky, and Michael Wotatzki. 2017. Measuring the reliability of hate speech annotations: The case of the european refugee crisis. *arXiv preprint arXiv:1701.08118*.
- Anna Schmidt and Michael Wiegand. 2017. A survey on hate speech detection using natural language processing. *SocialNLP 2017* page 1.
- Duyu Tang, Bing Qin, and Ting Liu. 2015. Document modeling with gated recurrent neural network for sentiment classification. In *EMNLP*. pages 1422–1432.
- Cynthia Van Hee, Els Lefever, Ben Verhoeven, Julie Mennes, Bart Desmet, Guy De Pauw, Walter Daelemans, and Véronique Hoste. 2015. Detection and fine-grained classification of cyberbullying events. In *International Conference Recent Advances in Natural Language Processing (RANLP)*. pages 672–680.
- William Warner and Julia Hirschberg. 2012. Detecting hate speech on the world wide web. In *Proceedings of the Second Workshop on Language in Social Media*. Association for Computational Linguistics, pages 19–26.
- Zeerak Waseem. 2016. Are you a racist or am i seeing things? annotator influence on hate speech detection on twitter. In *Proceedings of the 1st Workshop on Natural Language Processing and Computational Social Science*. pages 138–142.
- Zeerak Waseem and Dirk Hovy. 2016. Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *Proceedings of NAACL-HLT*. pages 88–93.
- Ellery Wulczyn, Nithum Thain, and Lucas Dixon. 2016. Ex machina: Personal attacks seen at scale. *arXiv preprint arXiv:1610.08914*.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of NAACL-HLT*. pages 1480–1489.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*. pages 649–657.