

hBert + BiasCorp - Fighting Racism on the Web

Olawale Onabola¹, Zhuang Ma²

Yang Xie³, Benjamin Akera¹, Abdulrahman Ibraheem¹,
Jia Xue⁴, Dianbo Liu¹, and Yoshua Bengio^{1,5}

¹Montreal Institute for Learning Algorithms (Mila), ²Carnegie Mellon University, CMU,

³Independent Researcher, ⁴University of Toronto,

⁵CIFAR Program Co-director

walexig4great@gmail.com, liudianbo@gmail.com

yoshua.bengio@mila.quebec

Abstract

Subtle and overt racism is still present both in physical and online communities today and has impacted many lives in different segments of the society. In this short piece of work, we present how we're tackling this societal issue with Natural Language Processing. We are releasing BiasCorp¹², a dataset containing 139,090 comments and news segment from three specific sources - Fox News, BreitbartNews and YouTube. The first batch (45,000 manually annotated) is ready for publication. We are currently in the final phase of manually labelling the remaining dataset using Amazon Mechanical Turk. BERT has been used widely in several downstream tasks. In this work, we present hBERT, where we modify certain layers of the pretrained BERT model with the new Hopfield Layer. hBert generalizes well across different distributions with the added advantage of a reduced model complexity. We are also releasing a JavaScript library³ and a Chrome Extension Application⁴, to help developers make use of our trained model in web applications (say chat application) and for users to identify and report racially biased contents on the web respectively.

1 Introduction

The internet has evolved to become one of the main sources of textual information for many people. Through social media, reviews, and comment sections across the internet, people are continuously consuming information through text.

With this, racially biased content has become more entrenched within the language of the internet. Racially biased content in this context refers to the attitudes or stereotypes expressed against marginalized races. This is often as a result of implicit bias resulting into hate speech. In this work, we attempt to automatically detect this racially biased content from data collected from the web, including comments from online news outlets such as Fox News and and comments from YouTube videos. We label this dataset with pointers to racial bias and use machine learning techniques to automate this task. Specifically, we implement BERT as a base model to do this. We also implement a browser extension as a tool to help people identify racially biased content in the information they are consuming. We will also be releasing our curated dataset - BiasCorp to allow more research to be done in this direction.

2 Related works

One of the earliest papers to investigate machine learning approaches for the automatic detection of racially-biased online content is (Greevy and Smeaton, 2004). The paper identified the potential use of bag-of-words, n-grams, and distributions of parts-of-speech tags as features for the task. Their bag-of-words features are informed by ideas from the field of information retrieval, and involve either word frequencies or counts of word occurrences. Using an SVM classifier, for bag-of-words features, they found that the use of frequency of words, rather than number of occurrence of words, yielded greater classification accuracies. The n-grams and parts-of-speech tags techniques were unavailable as of the time of their writing.

In (Warner and Hirschberg, 2012), authors

¹corresponds to:dianbo.liu/ yoshua.bengio@mila.quebec

²<https://doi.org/10.7910/DVN/KPBRLC>

³<https://github.com/walexig/biasjs>

⁴<https://github.com/walexig/biasjs-chrome-extension>

followed the definition of (Nockleby and John, 2000) by defining hate speech 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.” Their work focused more on detecting anti-Semitic hate speech. For their work, they created a dataset containing hate speech obtained from Yahoo! and the American Jewish Congress. Following the work of (Yarowsky, 1994), they employed hand-crafted template-based features. Apart from the fact that these features are hand-engineered, a potential drawback is their sheer size: a total of 3,537 features, which is prone to the curse of dimensionality. A counter-intuitive result reported by the paper is that the uni-gram features contributed best to classification accuracies. They used linear-kernel SVMs for classification.

The work of (C. et al., 2016) dealt with the broad category of abusive language. Authors of the work gave definitions for distinguishing between three categories of abusive language: hate speech which subsumes racial bias, derogatory remarks and profanity. Further, they described reasons why automatic detection of abusive language, which subsumes racial bias, is difficult. Reasons include: clever evasion of detection engines by users via the use of mischievous permutations of words (e.g. Niggah written as Ni99ah); evolution of ethnic slurs with time; role of cultural context in the perception and interpretation of slurs, as a phrase that is considered derogative in one culture might be perfectly neutral in another culture. Towards building their classification model, they employed four categories of features namely, n-grams, lexical features, syntactic/parser features, and word-level as well as comment-level embeddings. They found that character-level n-grams gave the highest contribution to the model’s accuracy.

The authors of (Burnap and Williams, 2016) also developed techniques for detecting multiple hate speech categories including the racially-based category. Towards creating their datasets, they harnessed hate speech event-triggers. For example, to create their racial bias dataset, they collected tweets in a two-week interval following the re-election of Barack Obama as U.S president. They explored a number of potential features towards building their classification algorithm: bag of words, lexicon of hateful terms, and typed

dependencies. In addition, they experimented into classification via SVMs versus classification via random forests, and reported that the former yielded superior performance over the latter. Also, they compared the use of classifiers trained for each hate speech category against the use of a single classifier trained on data spanning all categories. As expected, the specialized classifiers outperformed their multi-category counterpart. (Hasanuzzaman et al., 2017) followed the definition of (Gelber and Stone, 2007), which states that hate speech is: ”speech or expression which is capable of instilling or inciting hatred of, or prejudice towards, a person or group of people on a specified ground, including race, nationality, ethnicity, country of origin, ethno-religious identity, religion, sexuality, gender identity or gender.” The main research thrust of their work was to apply demographic embeddings (Bamman et al., 2014) , (Hovy, 2015), for the task of racial bias detection in tweets. Compared to other works such as (Burnap and Williams, 2016), for instance, a particularly distinguishing result of (Hasanuzzaman et al., 2017) is how their data extraction procedure is able to arrive at a better balanced ratio of racially-biased to non-racially-biased comments. For example, in the work, 40.58 percent of Canadian tweets were judged racially-biased by human annotators, whereas in (Burnap and Williams, 2016) only about 3.73 percent of the comments in the dataset are racially biased. Classification results using an SVM classifier revealed benefits of their proposed demographic embeddings over traditional features and embeddings. In (Saleh et al., 2020), the authors explored the detection of hate speech in White supremacist forums. They explored BiLSTM, logistic regression and BERT for their task. Also, they compared the use of domain-agnostic pre-trained word embedding (such as GloVe.6B.300d) versus the use of a domain-aware 300-dimensional word2vec embedding trained on the specific dataset used in the work. Results showed that BERT yields better results than both logistic regression and BiLSTM. Further, results proved the domain-aware embeddings to be superior to the pre-trained embeddings.

3 Method

3.1 Data curation and processing

The datasets used for training were obtained from discussion channels of online news media

by programmed web crawler based on Scrapy framework with all crawled data stored in PostgreSQL database. Since existing comments of online article were generally loaded by asynchronous API accessed by a specific key hidden in the articles before presenting them on website, the web crawler parsed keys for each article after completing a list with URLs of all articles waiting to be further crawled and then matched the keys with their corresponding API to retrieved stored comments for each article.

First, sentences containing neural racial words from a curated list were selected. Second, the sentiment score of each comment was calculated according to two lookup tables: a combined and augmented (Jockers, 2015) and Rinker’s augmented Hu and Liu (Tyler Rinker, 2016) (Hu and Liu, 2004) positive/negative word list as sentiment lookup values, and a racial-related English lookup table from Hatebase⁵. To guarantee these two tables influence the sentiment score consistently, the lookup values of the Hatebase table were adjusted by percentage. Then we extracted the data with bottom 20 percent of the sentiment score, and matched them up with other randomly selected comments appearing under the same articles or videos as random control. Finally, equal numbers of random controls are added into the data set, to ensure that approximately half of the data is racially discriminatory.

3.2 Model Architecture

Attention-based Transformer network (Vaswani et al., 2017) has been used widely across different natural language processing tasks. Based on the previous successes of the transformer network, we decided to use the BERT Architecture (Devlin et al., 2019) as our base model. Unlike previous variant of the attention-based language models such as (Radford et al., 2018), BERT learns to jointly conditions on the right and left context of the input representation at all the layers by randomly masking out segments of the input token. This is particularly useful for extracting contextual information from the input representation, and it’s very applicable to our use case. We aim to build a variant of the model that can generalize *sufficiently* well across different data distributions⁶. The notion of *sufficiency* is evaluated by training,

⁵<https://hatebase.org/>

⁶distributions here implies different use cases or data environments/sources

validating and testing our model on data across the different sources. We fine-tune the pretrained BERT model on our curated dataset rather than training from scratch (this choice was based on empirical results). We are releasing a JavaScript library for developers to use our pretrained model in front facing applications such as chat app, to flag down racially biased comments. Consequently, we need to optimize for the model complexity without sacrificing performance gain. BERT has a huge number of parameters / large model size. Other methods have been employed to reduce the complexity without hurting the performance, such as knowledge distillation (Sanh et al., 2019) and quantization (Zafir et al., 2019). It has also been proven that pruning the weights of the pretrained model do not necessarily affect the model performance, within acceptable ‘thresholds’ (Gordon et al., 2020). In a similar fashion, we aim to reduce the complexity of BERT without sacrificing performance by replacing certain layers with the Hopfield layer (Ramsauer et al., 2020). Hopfield layer can be used to replace the attention-based layer of the BERT model; as it has been shown to approximate the functionality of the attention mechanism with a new Energy update rule (modified version of the Hopfield network extended to continuous state representation). The learning dynamics of BERT as shown in (Ramsauer et al., 2020) shows that the attention heads in the higher layers are mostly responsible for extracting task-specific features from the input representation. We replaced the self-attention mechanism in the last X layers of the pretrained BERT model with a Hopfield layer, where X is an hyperparameter. In a similar approach described in (Vaswani et al., 2017), we use residual connection around the Hopfield sub-layer, followed by layer normalization (Ba et al., 2016). It has been shown that residual connections help propagate positional information across layers. The replaced Hopfield layer drastically reduced the parameter size of our model. To further improve the performance of the model, we use the Hopfield Pooling layer which acts as both a permutation equivariant layer and pools generated embedding from the modified BERT model. The Hopfield pooling layer also acts as a form of memory to store the hidden state of the last layer in the modified BERT model. Finally, we add a classification layer on top of the pooling layer for the task in question.

3.3 Model Training

Given the disparity between the annotators for each sample in our dataset, averaging the labels with the confidence scores as weights might be noisy. We computed the coefficient of variation CV among annotators for each sample in our dataset. Using the recommended (JUDICE et al., 1999) (Veit et al., 2017) CV of 0.2 for the bias scores would imply dropping 90% of the dataset as seen in 2. In order to fully utilize the dataset and effectively manage the disparity between the annotators, we formulate a loss function \mathcal{L}_{model} given by

$$L_{model} = 1/N \sum_{i=1}^N CE(p(x_i), q(x_i)) \quad (1)$$

where $CE(p(x_i), q(x_i))$ is the cross entropy between $p(x_i)$ and $q(x_i)$ for the i th sample, and N is the size of the dataset.

$$CE(p, q) = - \sum_{i=1}^c p_c(x) \log(\epsilon + q_c(x)) \quad (2)$$

$q_c(x)$ is the predicted probability of sample x in class c , equivalently, the output probabilities from the model and ϵ is for numerical stability. $p_c(x)$ is the probability of sample x in class c , equivalently, $p_c(x)$ is a $c - length$ vector with entries such that $\sum_{i=1}^c p_c(x) = 1$. The entries of $p_c(x)$ are the normalized confidence scores of the annotators with indices given by the respective voted classes. As an example, following the algorithm described in 1, for a given sample shown in figure 1; the bias scores of the 3 different annotators with their confidence level is represented with an array of tuples, X where each tuple, (b_i, s_i) is the bias score b_i with the associated confidence score, s_i by annotator i . To calculate $p_c(x)$, we first normalize the confidence scores across the 3 different annotators such that $\sum_{i=1}^3 s_i = 1$. The resulting $p_c(x)$ for the entry, S , shown in 1 is

$$X = \left[(4, 4), (3, 3), (2, 5) \right]$$

$$X_{norm} = \left[(4, 0.3333), (3, 0.25), (2, 0.4167) \right]$$

$$p_c(X) = [0., 0., 0.4167, 0.25, 0.3333, 0.]$$

3.4 Evaluation Task and Metrics

We evaluate the model performance across the validation and test set, given that they are from

Algorithm 1: Compute $p_c(x)$ for a sample

x

Result: $p_c(x)$

Input: An array of target scores t , and array of confidence scores s where $s[i]$ is the confidence score by **annotator** i for choosing target score $t[i]$

Both arrays are of equal length N where N is the number of annotators. C is the number of classes (equivalently the range/max of possible target scores if scores are integer.)

Step 1: Initialize

$p_c \leftarrow [.0 \text{ for } _ \text{ in } C]$

Step 2: Calculate normalizing constant K

$K \leftarrow \sum_{i=1}^N s_i$;

Step 3: Set the values of p_c

for i in N **do**

$class_index \leftarrow t[i]$;
 $p_c[class_index] \leftarrow \frac{s[i]}{K}$;

end

different distributions or sources. The test set contains only comments from YouTube while the validation set was randomly sampled from Fox News and BreitbartNews. The particular choices were due to the fact that the first batch of the dataset used for training contained very relatively few samples from YouTube. We evaluate our approach using two methods; multiclass classification and multiclass-multilabel classification.

Using the multiclass approach, for a given sample, k and using the method described previously in calculating the target class, the class with the maximum confidence score was used as the target. We calculate the average precision for each class, AP_c and the mean average precision MAP averaged over the entire dataset with size N along the class dimension d as described in (Veit et al., 2017)

$$AP_c = \frac{\sum_{k=1}^N Precision(k, c) \cdot rel(k, c)}{number\ of\ positives} \quad (3)$$

$$MAP = 1/d \sum_{c=1}^d AP_c \quad (4)$$

	text	bias_1	bias_2	bias_3	confidence_1	confidence_2	confidence_3
0	" Democrats needed someone like Obama, half w...	4	3	2	4	3	5

Figure 1: Sample annotation

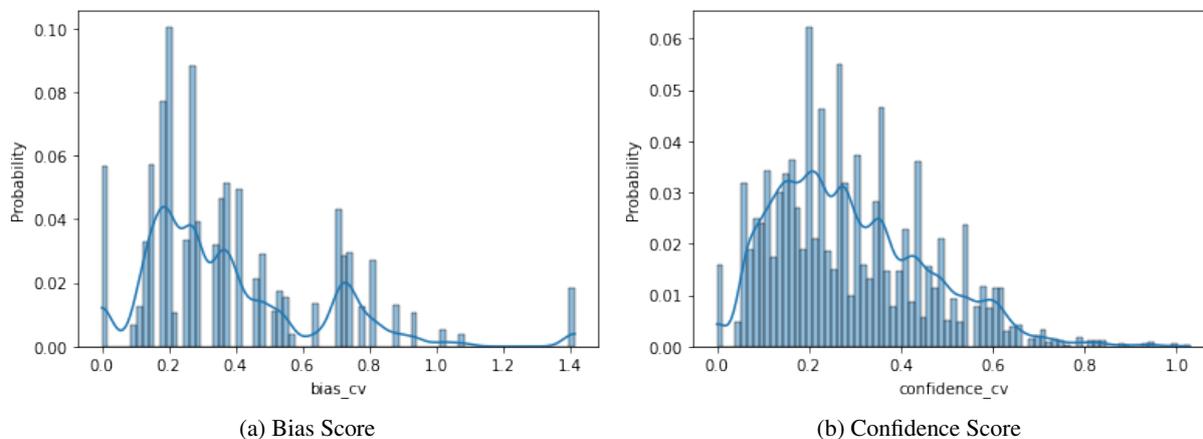
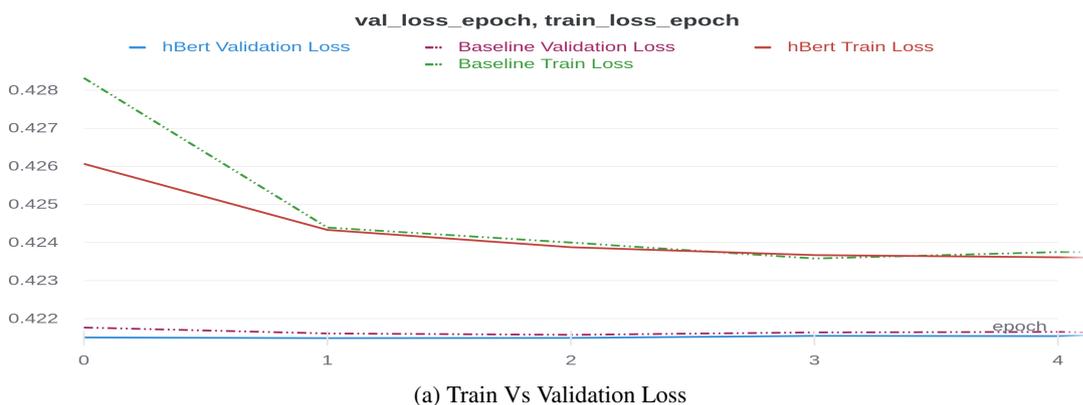


Figure 2: Confidence of Variation



(a) Train Vs Validation Loss

Model	TopK Accuracy			mAP	F1 @ k			IoU @ k		
	1	2	3		1	2	3	1	2	3
Baseline	0.6015625	0.703125	0.7890625	0.29355	0.5859	0.6953	0.7734	0.2102	0.2114	0.2102
hBert	0.640625	0.703125	0.765625	0.3501	0.6562	0.7109	0.8125	0.2266	0.2165	0.2281

Table 1: Test Metrics for selected trial for each model configuration

Model	AP					
	0	1	2	3	4	5
Baseline	0.2205	0.0967	0.1344	0.9564	0.1103	0.2340
hBert	0.1195	0.1111	0.2132	0.9607	0.5049	0.1914

Table 2: The Average Precision (AP) for the different classes

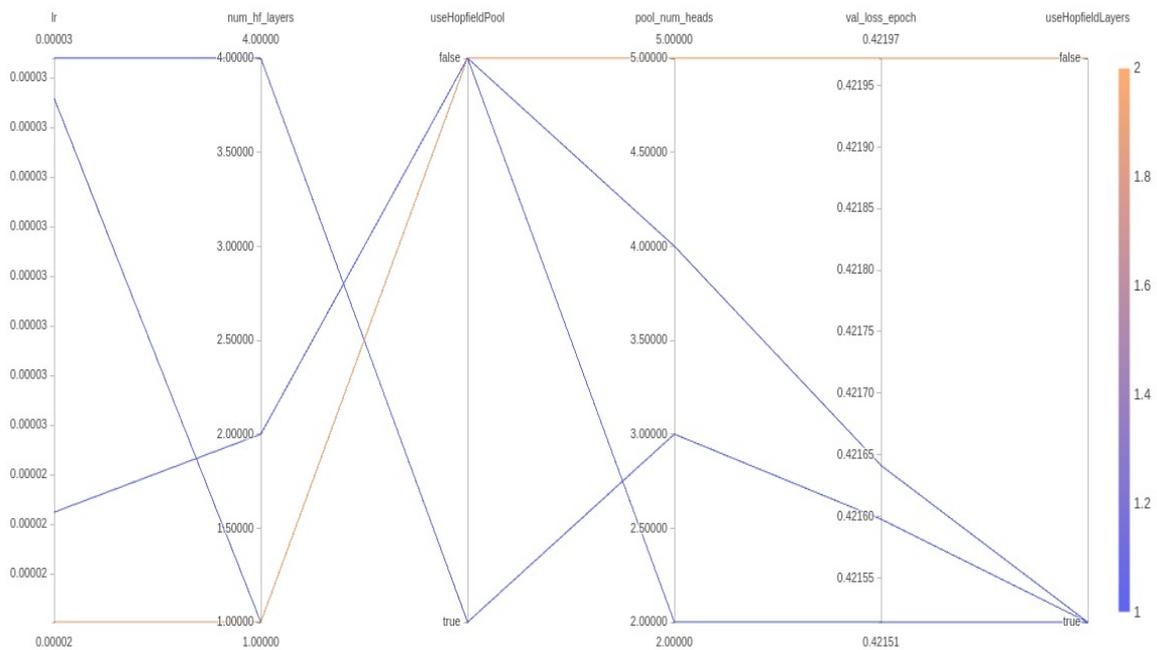


Figure 3: Parallel Coordinate Graph for multiple runs/trials across model configurations

The model configuration is the Baseline when the target variable (*useHopfieldLayers* in the graph is False.

The *useHopfieldPool* variable denotes whether the Hopfield Pooling layer was used

The *lr*, *pool_num_heads*, *num_hf_layers*, *val_loss_epoch* variables in the graph are the learning rate, the number of heads in the Hopfield Pooling Layer (if used), the number of Hopfield Layer and the validation loss respectively

. With a reduced model complexity, the hBert performs relatively as good as the baseline

where $Precision(k, c)$ is the precision for class c for the k th sample and $rel(k, c)$ is an indicator function that is 1 if the predicted and the target class for sample k is positive. We also report the $topK$ accuracy, for $k = [1, 3]$ since we had a max of 3 annotators for each k .

Using the multilabel approach, for a given sample, k and using the method described previously in calculating the target class, we take the top k classes as the target classes. We do the same for the predictions (obtained after passing the output logits through a softmax function). We compute the AP_c (for each class), mAP , $F1$ score, and IoU

4 Experiments

4.1 Training Details & Result

We run a multi-objective hyperparameter search (using Optuna(Akiba et al., 2019)) optimizing for the following parameters: validation loss, FLOPs (indicative of the model complexity and ultimately the inference time), mAP on the validation and test set, and the Intersection over Union IoU scores (also known as the Jaccard Index) for the $topk$ for $k = [1, 3]$ transformations described above. We use 4 NVidia V100SXM2 (16G memory) GPUs on a single node, with batch size of 32. We reduced the batch size (instead of say 64) because we had to run multiple trials and to avoid the notorious OOM error. For each model configuration, we run 10 trials with 5 epochs each. As seen in 3, the hBert perform relatively better with a reduced model complexity. In 1, the models predictions were more accurate for an increasing k . The hBert perform better than the Baseline for the Top1 accuracy. The $F1$ scores and Jaccard Index (IoU) for the hBert were relatively higher for $k = [1, 3]$. The mAP , which is the average of the AP_c over the classes, is relatively low because of the low performing classes as seen in 2

4.2 Data statistics

The data set contains 139,090 rows, and 67.70 percent of their sentiment scores are negative. Their average sentiment score is -0.1422, and the median value is -0.1203, ranging from -3.6206 to 2.1414. 66,998 of them are comments from Fox News, with an average sentiment score of -0.0997 and a median of -0.0884, ranging from -2.8591 to 2.1414. And 63,948 of the data are comments from Breitbart News, with an average sentiment score of -0.1760

and a median of -0.1721, ranging from -3.6206 to 1.3576. And 8,144 of the data are comments from YouTube, with an average sentiment score of -0.2259 and a median of -0.2694, ranging from -3.3000 to 1.4673. In this work, we used the first batch of the dataset; which have been manually annotated using Amazon Mechanical Turk. After pre-processing the input text (removing irrelevant tokens such as mentions), the maximum length was 478 (it was 623 before preprocessing).

5 Discussion

In this work we have shown a way to detect racial bias in text. We experimented with a BERT-based model as we aim to reduce model complexity without sacrificing much of the performance. We also discussed the BiasCorp, a manually labelled dataset containing racially biased comments from Fox News, BreitbartNews and YouTube. To enable developers make use of our pretrained hBERT model, we are releasing a Javascript Library, optimized for inference on the edge. A Chrome Extension will also be available for users to help report and identify racially bias text on the web. We also plan to extend this work to other forms of biases such as Gender. In a future work, we plan to further reduce the model complexity by using Gaussian Kernel as described in (Ramsauer et al., 2020) and other quantization tricks.

Acknowledgments

This research was enabled in part by support provided by Calcul Québec (www.calculquebec.ca) and Compute Canada (www.computeCanada.ca)

References

- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#).
- David Bamman, Chris Dyer, and Noah A. Smith. 2014. Distributed representations of geographically situated language. In *Proceedings of the Association for Computational Linguistics*, pages 828–834.
- Pete Burnap and Mathew Williams. 2016. Us and them: identifying cyber hate on twitter across

- multiple protected characteristics. *EPJ Data Science*, 5:1817–1853.
- Nobata C., Tetreault J., Thomas A., Y. Mehdad, and Chang Y. 2016. Abusive language detection in online user content. In *Proceedings of the 25th international conference on world wide web*, pages 145–153.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Katharine Gelber and Adrienne Sarah Ackary Stone. 2007. *Hate Speech and Freedom of Speech in Australia*. Federation Press.
- Mitchell A. Gordon, Kevin Duh, and Nicholas Andrews. 2020. [Compressing bert: Studying the effects of weight pruning on transfer learning](#).
- Edel Greevy and Alan Smeaton. 2004. Text categorisation of racist texts using a support vector machine. *JADT 2004 : 7es Journées internationales d'Analyse statistique des Données Textuelles*, pages 533–544.
- Mohammed Hasanuzzaman, Gael Dias, and Andy Way. 2017. Demographic word embeddings for racism detection on twitter. In *Proceedings of the 8th International Joint Conference on Natural Language Processing*, pages 926–936.
- Dirk Hovy. 2015. Demographic factors improve classification performance. In *Proceedings of the Association for Computational Linguistics*, pages 752–762.
- Minqing Hu and Bing Liu. 2004. [Mining and summarizing customer reviews](#). pages 168–177.
- Matthew L. Jockers. 2015. *Syuzhet: Extract Sentiment and Plot Arcs from Text*.
- MARCELO JUDICE, Muniz Augusto, and Roberto Carvalheiro. 1999. Avaliação do coeficiente de variação na experimentação com suínos. *Ciência e Agrotecnologia*, 23.
- Nockleby and John. 2000. Hate speech. In *Encyclopedia of the American Constitution*, pages 1277–1279.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding with unsupervised learning.
- Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Thomas Adler, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Kjetil Sandve, Victor Greiff, David Kreil, Michael Kopp, Günter Klambauer, Johannes Brandstetter, and Sepp Hochreiter. 2020. [Hopfield networks is all you need](#).
- Hind Saleh, Areej Alhothali, and Kawthar Moria. 2020. Detecting white supremacist hate speech using domain specific word embedding with deep learning and bert. In *Submission to Information Processing and Management*, pages 533–544.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *CoRR*, abs/1910.01108.
- Vitalie Spinu Tyler Rinker. 2016. [sentimentr: Dictionary based sentiment analysis that considers valence shifters](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#).
- Andreas Veit, Neil Alldrin, Gal Chechik, Ivan Krasin, Abhinav Gupta, and Serge J. Belongie. 2017. [Learning from noisy large-scale datasets with minimal supervision](#). *CoRR*, abs/1701.01619.
- J. Warner and J. Hirschberg. 2012. Detecting hate speech on the world wide web. In *Proceedings of the Second Workshop on Language in Social Media*, pages 19–26.
- David Yarowsky. 1994. Decision lists for lexical ambiguity resolution: Application to accent restoration in spanish and french. In *Proceedings of the ACL*, pages 88–95.
- Ofir Zafrir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. 2019. [Q8BERT: quantized 8bit BERT](#). *CoRR*, abs/1910.06188.