

A Comprehensive View of the Biases of Toxicity and Sentiment Analysis Methods Towards Utterances with African American English Expressions

Guilherme Andrade, Luiz Nery,
Fabrício Benevenuto, Flavio Figueiredo
UFMG
guilherme.hra,luiznery,fabricio,flavio@dcc.ufmg.br

Savvas Zannettou
TU Delft
s.zannettou@tudelft.nl

ABSTRACT

Language is a dynamic aspect of our culture that changes when expressed in different technologies and/or communities. On the Internet, social networks have enabled the diffusion and evolution of different dialects, including African American English (AAE). However, this increased usage of different dialects is not without barriers. One particular barrier, the focus of this paper, is on how sentiment (Vader, TextBlob, and Flair) and toxicity (Google’s Perspective and models from the open-source Detoxify) scoring methods present biases towards utterances with AAE expressions. In particular, AI tools cannot understand the re-appropriation of the terms, leading to false positive scores and biases. Here, we study the bias of Toxicity and Sentiment Analysis models based on experiments performed on Web-and spoken English datasets.

KEYWORDS

African American English, AAE, Bias, Toxicity, Sentiment

1 INTRODUCTION

In recent decades, we have witnessed a substantial rise in Internet usage. According to [52], Internet users increased from approximately 400 million in 2000 to 4.7 billion in 2020. With this increase in usage, it is natural that Web applications enable a wide diversity of social groups to interact among themselves and with other groups. Since such applications foster a more open and dynamic form of speech, a natural increase in the written form of dialects that previously were predominantly seen in the spoken form [9] occurred. However, such massive amounts of textual data make manual content moderation impracticable. In other words, the heavy usage of social media has evidenced the urge for automatic moderation tools that measure and moderate improper behavior online. One of the main concerns is the public display of negative/toxic sentiments against a person or specific group, more drastically when the target is a minority group historically marked with discrimination and stereotypes. The necessity of dealing with the increasing number of deviating content has led many researchers and companies to use AI tools to identify such events [51].

Concurrently to the increase in Web usage, African-American English (AAE) has gone from being seen as a marginalized dialect of English to a consolidated vernacular of the language [28]. Like most

dialects, the AAE was initially heavily used in spoken form and had the Web as a crucial influence on its emergence in the written form [9]. However, the Web is not only a disseminator of cultural aspects of our society but also a vehicle where toxicity campaigns against African Americans are prone to occur. Even though several websites have well-defined community guidelines, user anonymity and lack of unaccountability leave room for misbehavior.

The aforementioned rise in AI moderation tools (such as Google’s Perspective [37] and others [23, 51, 66]) aim to reduce the amount of negative or toxic utterances online. Overall, such tools rely on Machine Learning (ML) models that help determine proper and improper utterances. Nevertheless, as previous research has discussed, automatic content moderation can backfire and present biases towards minorities [11, 34, 53, 60]. For instance, a tool for toxicity analysis may present high scores for non-toxic AAE sentences for no apparent reason. We show examples of toxicity and sentiment analysis models employed in online text to depict this issue. We point out that it is quite easy to find problematic utterances when using slang terms such as “n****s”. In Table 1, we contrast three pairs of sentences based on their toxicity/negative sentiment levels.

Why does the problem arise? From a linguistic perspective, dialects may inherently manifest behaviors and cultural aspects of the groups in which they were created [5, 19–21]. Terms such as “n****r” are problematic for AI models since both the term and its variations have a historical pejorative usage¹. Nevertheless, this same term was re-appropriated by the black community, so its use ceased to be considered problematic when used by people inside the black community. Suppose such a fine line between causal speaking and offensive discourse is problematic from a human and computational perspective. In that case, these interpretations are confounding to automatic content moderation tools. In other words, toxicity/sentiment analysis tools are usually developed using manual rules or supervised ML techniques that employ human-labeled data to extract patterns. The disparate treatment embodied by machine learning models usually replicates discrimination patterns historically practiced by humans when interacting with processes in the real world. Due to biases in this process, a lack of context leads models to a concerning scenario where minorities do not receive equal treatment [1, 14, 26, 58].

This discussion leads to the research question behind our paper: *Is there a systematic bias on toxicity/sentiment analysis towards AAE?* To better understand this issue, we present a broad-scale analysis. To do so, we manually curate a dataset of African American English Expressions [54, 63] (these sources were recommended by

In: Proceedings of the Brazilian Symposium on Multimedia and the Web (WebMedia’2024). Juiz de Fora, Brazil. Porto Alegre: Brazilian Computer Society, 2024.
© 2024 SBC – Brazilian Computing Society.
ISSN 2966-2753

¹<https://en.wikipedia.org/wiki/Nigga>

	Score	All my friends on the porch and never in the house	All my n****s on the porch and never in a house	You're white	You're black	I can't forget you	Cant fuhgit you
Persp. (ML)	[👍=0, 👎=1]	0.2396	0.7886	0.2546	0.4256	0.0406	0.2359
Detox (ML)	[👍=0, 👎=1]	0.0012	0.6145	0.8766	0.9718	0.0257	0.7601
Detox U (ML)	[👍=0, 👎=1]	0.0013	0.6842	0.2549	0.8903	0.0162	0.5332
Vader (Lexic)	[👍=0, 👎=1]	0	0	0	0	0	0
Textblob (Lexic)	[👎=-1, 👍=1]	0	0	0	-0.1666	0	0
Flair (ML)	[👎=-1, 👍=1]	0.9830	-0.7296	0.9994	0.9992	0.9957	0.9111

Table 1: Examples of Toxicity Scoring and well as the Sentiment Scoring. Good scores (👍) are colored green, bad scores (👎) are colored red. Other shades represent scores in between. For Textblob and Flair, 👎=-1 are utterances of negative sentiment. In Perspective, Detoxify, and Vader, a negative sentiment or toxic utterance has a score of 👎=1.

the organizers of the to-be-released Oxford Dictionary of African American English²). These expressions complement four different utterance datasets with some demographic information on race (i.e., interviews from African American individuals, AAE utterance vs non-AAE utterance labels, and author-supplied labels). *We emphasize that we cannot state how a speaker identifies regarding her/his race for some datasets. AAE may also be employed by non-African Americans. However, we interpret our results using AAE expressions from our expressions dataset (made available).*

The models we study can be divided into toxicity (Google’s Perspective [37], Detoxify, and Detoxify Unbiased [30]), and sentiment analysis (Flair [3], TextBlob [42], and Vader [35]) models, but also can be segmented into machine learning-based (Google’s Perspective, Flair, Detoxify, and Detoxify Unbiased), and lexical, or rule, based (Vader, and TextBlob) models. Our main contributions are:

- (1) We present a broad-scale analysis of biases toward utterances with AAE expressions in six out-of-the-shelf models;
- (2) To do so, we focus on unveiling if there is a systematic tendency for AAE utterances to be deemed more toxic or negative sentiment by several models of datasets of different natures (tweets, closed captions, and spoken interviews). To reach our results, we manually transcribed a dictionary of AAE expressions and used the number of such expressions in an utterance as an explanatory feature;
- (3) Other control features include Lexical Analysis [49], and grammar-based Part of Speech Tagging (PoS) labels for words in utterances. Overall, we discuss which characteristics of the utterance lead the model to deem it as toxic or of negative sentiment. The number of AAE expressions is a recurrent statistically significant feature;
- (4) Using recent language models [62], we show that in some datasets, even utterances from African-American (AA) speakers that have a similar meaning to those from non-AA speakers, models will, in several cases, score the sentence from non-Whites with more toxic/more negative scores.

Before continuing, we point out that our work is not the first to study the biases of similar models towards minorities [1, 14, 26, 58]. However, we complement prior endeavors with broader-scale analysis. Previous methods focus on a single dataset or model and do

not perform the lexical and grammar-based analysis we do here. We further point out that our goal in this paper is not to pinpoint models with the best accuracy. We focus on showing AAEs and comparing if there is a tendency across several models. The datasets where we show this issue range from online texts from Twitter [9, 10], spoken English datasets gathered by linguists [38, 50], and online single speaker English from YouTube movie reviews. The YouTube dataset (see Section 3) was a manual effort toward gathering data with fewer confounding factors (i.e., single-speaker videos). This dataset is made available to the community to improve the current and yet-to-come NLP models³.

Our results show that biases are more prominent on online datasets, such as Twitter and YouTube, and less strongly but still present in spoken English interviews. Our research shows that using AAE expressions will likely lead to sentences being deemed more toxic, even when sentences are similar to those with non-AAE expressions. Overall, system developers may use these findings to determine what model type shall be employed (sentiment analysis vs. toxicity scoring) or whether ML vs. lexical-based models are more adequate for their application. More importantly, our findings show that even considering “unbiased models” [30], ML models still present a bias towards utterances with AAE expressions. Indicating that AAE speakers may still face unwarranted moderation online.

2 BACKGROUND AND RELATED WORK

This section presents an overview of the literature. We begin discussing the sociolinguistics of English online, as the increased usage of AAE expressions online is a primary motivation behind our work. Subsequently, we also discuss the motivation behind sentiment analysis tools, available alternatives, their major strengths and shortcomings, and how toxicity relates to sentiment analysis. Next, we discuss bias in machine learning methods and how they can negatively influence individuals online and suppress the discourse of minority groups. Finally, we focus on those papers most related to ours and present a statement on the novelty of our research.

²<https://hutchinscenter.fas.harvard.edu/odaae>

³https://anonymous.4open.science/r/aae_bias-D396/data/aae_terms_black_talk.yaml

2.1 Sociolinguistics of AAE Online

As a research field, sociolinguistics focuses on studying how social context affects the usage and evolution of language. Overall, humans take part in several speech communities throughout their lives [21], and even the same human being may communicate in different variations of English depending on the community she/he is interacting with. With the rise of the Social Web over the 2000's and 2010's, the field also focused on how Web communities affect language [20, 21]. In particular, Friedrich and Figueiredo [21] argue that hundreds of years after the invention of the printing press, the written usage of English as a language appeared to be "evolving" to a standard or uniform English. However, with the Web, community and individualized language use has increased over recent years. AAE is an example of such a case [9, 20, 21], where the dialect has experienced a rise in usage (particularly online) in recent years.

One example of the expansion of AAE in recent years comes from the movement known as *signifyin'* [19]. In other words, when *signifyin'* one expresses their race via particular dialects, such as AAE, on social media. This expression is utilized to resist the oppression present in one's day lives [19, 46]. Regarding how AAE is spread online, some authors argue that the dialect spreads initially from Web users from large cities to smaller communities in wave-like, or viral, patterns [17]. Frierich and Figueiredo state that:

"With the Internet, we have witnessed a change in this scenario. Gender and racial/ethnic activism have become quite strong online and have served not only to spread the debates but also to add new layers to them – such as the complex construction of identities in cyberspace. And again, we must say, English has been quite present in this new picture, mainly because of its lingua franca status."

For several decades, the AAE dialect has also been studied offline [63]. Over the years, some attempts have been made to catalog different expressions from the dialect [8, 15, 54]. Oxford and Harvard are also organizing a dictionary called the Oxford Dictionary of African American English (ODAAE).

Given that language constantly evolves, we aimed to collect a recent corpus of AAE expressions. To do so, we contacted the organizers of the ODAAE, asking if they could share the list of expressions used in their dictionary. The organizers kindly denied it because ODAAE is still a work in progress. Nevertheless, they did suggest that we use expressions from Smitherman's Black Talk Dictionary [54] as it is a large and somewhat recent corpus. In our research, we manually transcribed every expression from this dictionary as our AAE expression list.

2.2 Sentiment Analysis and Toxicity Models

Sentiment Analysis identifies sentiments and quantifies their intensity (positive or negative sentiment) in utterances. Current Sentiment Analysis models may be classified into two major categories, namely, machine learning-based (ML models) and lexical-based approaches, described below. These methods have been employed since the mid-2000s, and one of the major motivations behind building such approaches was to rate user reviews online.

ML-based Sentiment Analysis models [3, 24, 55, 61, 67] are built over a sample of data points comprising as many examples as possible from positive and negative sentences. Usually, the learning

procedure targets data drawn from a context of interest (e.g., Twitter, Facebook, Marketplaces, etc.), and humans manually label this dataset to train models. This family of methods often benefits from complex word representations and can grasp deeper relationships implied in daily conversations. Lexical-based approaches [33, 35, 49, 57], on the other hand, begin by listing seed words considered to be representatives of groups of emotions. Once the seed list is complete, it is incremented with similar words and synonyms. Such approaches must actively deal with normal word usage that may change the intent/intensity of the sentence, such as negations, punctuations, capitalization, etc. Since this approach is based on the human understanding/application of terms and expressions, their performance on novel datasets may have less statistical variance.

Compared to lexical-based methods, the ML models rely on the vector representation of terms and utterances [44]. Such vector representations are used as inputs to train supervised methods. Collecting high-quality labels to train such models is a difficult-to-reach pre-requisite (discussed below). On the other hand, lexical-based approaches need to explicitly address negations, punctuations, out-of-vocabulary occurrences, and more complex relations between terms [51]. To address the gaps left by each family of methods, authors have also proposed hybrid solutions [47, 59, 64].

More recently, we have seen a rise in Toxicity Classification (compared to Sentiment Analysis) models [30, 37, 40]. Most, if not all, of these approaches are ML-based. Toxic speech is usually considered to be an umbrella term that comprises hate speech, abusive language, racism, and so on [31, 41]. Despite the efforts to address toxic speech, there is not a clear agreement about what it means for a sentence to be toxic. Dixon et al [16] defines toxicity as rude, disrespectful, or unreasonable language.

Due to the lack of consensus on toxicity, the inherent ambiguity of labeling sentences presents an issue to ML models. The vast majority of datasets use human labelers, which are influenced by their previous experiences and, most of the time, do not have access to the underlying context from which the respective sentence was drawn. This subjectivity and lack of context may cause considerable labeling issues. For example, Kumar et al. [41] state that people who have suffered harassment in the past are more prone to label random sentences from some social networks as toxic than those who did not face such problems. Maybe due to its less restrictive definition and to the capacity to encompass many types of harassment, toxicity models are actively used in practice to moderate discourse in many platforms [2, 43, 48]; however, with some known bias problems.

2.3 Biases Towards Minorities Online

We now discuss prior work on the biases of Web datasets and AI models. Starting from Jia et al. [36], the authors investigated the proportions in which men and women appeared in news articles' images. The authors found that men are considerably more frequently represented than women. Garcia et al. [22] also described a consistent bias towards men in Twitter content. On Twitter, female users tend to describe more events in which men play important roles. Babaeianjelodar et al. [4] explored the nuances of gender biases over ML models. In all datasets considered, models perform disparately against unprivileged subgroups. Similar findings were raised by several other authors considering countries [27], age [14],

religion [1], and sexual orientation [16]. Regarding dialects, Blodgett et al. [9] studied how language characteristics can change considerably within the same country. The work focuses on learning distinguishable features between Mainstream American English (MAE) and AAE with a geographic context. In [25], the authors also present another clear differentiation between English focusing on Drag queens. Here, the authors find that transgender individuals have a speaking characteristic consistently seen as more toxic.

As Bamman et al. [7] states, language is always situated within a context. Neglecting this surrounding context leads to disparate treatments. For example, transgender individuals may be speaking a dialect deemed toxic if used by someone outside of the community. However, this may be a defense mechanism to cope with tough situations imposed by society [25]. Similar language signals are passive within the black community and the AAE dialect. Studies were already performed to comprehend and measure how much ML models are biased against AAE speakers [6, 9, 53].

Overall, we can state that nowadays, it is not hard to find discrimination episodes involving AI models [1, 14, 16, 34]. For example, Abid et al. [1] interacted with a conversational artificial intelligence model touching religion-related subjects and noting the inner associations with the topic. Finally, they found a consistent bias associating Muslims with terrorists (in 23% of the test cases) and the Jewish with money (in 5% of the test cases). In the opposite direction, mitigating such biases are also common [4, 9, 11, 16, 31]. Nevertheless, as studied by Gonen et al. [26], persistent bias may stick with the model even after active effort has been applied to remove it. We also observe this as we use unbiased versions of recent models. Since the complexity of the ML model has increased in the last few years, we could expect the bias to be more elaborate and complex to fight against. This leads us to the problem of using biased models for sensible tasks that may perpetrate harmful behavior.

2.4 Related Work and Research Novelty

We now detail prior work that is most related to ours (e.g., evaluating and unveiling biases in similar models) [13, 18, 29, 32, 39, 45, 53, 56, 65]. Before doing so, we initially point out that the studying the biases of NLP models towards race is a well-established research topic and the survey of Field et al. [18] presents a recent overview on this topic. The authors of this survey analyzed 79 different papers on race and NLP systems. Overall, the consensus is that NLP still encodes racial biases (something we also observe) and that race is commonly studied as a limited categorical variable. Here, we take a step towards a broader view of the issues by incorporating in our study (1) a novel list of AAE expressions, (2) grammatical features (PoS), and (3) linguistic features (LIWC) to understand biases.

One of the vanguard efforts of looking into the biases of toxic scoring APIs (Perspective in particular) was performed by Sap et al. [53]. In contrast, Kiritchenko et al. were among the first to study the bias of sentiment analysis models [39]. Starting with the former, the authors discuss how datasets are biased and how models propagate such biases. However, unlike our work, the authors only study Twitter datasets and do not present statistical analysis on how utterance features (grammar, linguistics, and usage of AAE expressions) may explain biases. The authors also only focus on Perspective as an out-of-the-shelf model. The former, focused on

sentiment analysis, compares over two hundred models from a well-established information retrieval sentiment classification challenge. However, the author does not use real-world datasets as we do and focuses their analysis on sentences of similar meaning but with small changes in words related to gender, occupation, and race. This approach is similar to the adversarial attacks described next.

Although not focused on measuring biases, the work of Hosseini et al. [32], and Gröndahl et al. [29] both show how small changes to a sentence will change model scores. We present a different view on this finding by incorporating semantic similarity using language models [62] and finding that even expressions that are semantically close to one another will differ in scores depending on the number of AAE expressions used. Similarly, Davidson et al. [13] discussed some challenges in differentiating hate speech from other offenses. This provides evidence of how language is nuanced, and models have problems with these small nuances. Wright et al. [65] provides a tool called RECAST that helps users pinpoint words that need changing in order to reduce the toxicity of a score.

Regarding human evaluation of models, we refer to the recent effort of Muralikumar et al. [45]. Here, the authors evaluate Perspective and contrast how human scores align with the model. Overall, the score from Perspective is a good predictor (based on a Logistic regression) of human labels (“toxic”, “hard to say”, “non-toxic”). However, agreement is not always present, with the model still making mistakes. The authors suggest that using a score cut-off of 0.55, i.e. if the model scores over this value classify the utterance as toxic, will make model outcomes agree with users 50%.

Testing different definitions of fairness is also an active field of research [12], with software tools being developed just for this task in NLP models [56]. We complement these efforts by showing that out-of-the-shelf models and API still require further testing.

Our research differs from previous works by investigating biases in models of different families (ML-based and lexical-based methods) and throughout four datasets representing different contexts (in-person conversations, single-speaker movie reviews, and personal social media posts). Here, we focus on out-of-the-shelf methods and those already applied in real-world forums. Unlike prior work, we employ different statistical features of utterances to show how the presence of AAE expressions will lead models to rate sentences as more toxic or of negative sentiment. Our statistical analysis also provides insights into what features models explore to reach their scores. Finally, the datasets we employed were collected to isolate confounding factors. That is: (1) we do not use any dataset used to train models or with toxicity/sentiment labels; (2) one of our datasets is focused on single-speaker movie reviews, controlling for discourse as a confounder; (3) we also compare models on well-established linguistic datasets on single-person interviews.

3 DATASETS AND PRE-PROCESSING

We explore four datasets of different natures to understand the extent of biases in toxicity/sentiment analysis models and when they present themselves more strongly. Initially, we use the Twitter AAE dataset [9, 10]. This dataset is interesting as it contains tweets classified as AAE or Mainstream American English (MAE). Tweets were classified using an ML model, and we consider a subset of tweets where the model predicted over 80% probability for each

Dataset	Demographic	# Documents	# Sentences	# AAE Expr.	AAE Expr. per Doc.
YouTube	AA Speaker *	150	17828	18308	122.05 (▲ 43%)
	non AA Speaker *	484	41464	42729	85.67
Twitter	AAE Tweet	250	250	372	1.49 (▲ 43%)
	MAE Teet	250	250	259	1.04
CORAAL	AA	142	64493	61651	434.16 (● 9%)
Buckeye	Caucasian	39	19304	18712	479.79

Table 2: Datasets statistics. * Indicates that two (agreeing) authors inferred the demographic. Still, our analysis does not use it as a variable (we rely on the # of AAE Expr.).

class. This is a well-established dataset for AAE utterances used by other endeavors [53]. Twitter is one of the major platforms where one would expect that toxicity and sentiment analysis models could mitigate unwanted behavior. On the negative side, as the dataset contains general Tweets, it does not control for confounding factors such as dialogues, debates, and potentially controversial topics. Thus, we complement this research with two other datasets.

Our YouTube dataset comprises subtitles extracted from YouTube movie reviews with a single speaker discoursing about a unique topic per video. The topics are movies from Rotten Tomatoes 100 Best Movies of All Time. We targeted single-speaker videos to control for any confounding variables that may appear with dialogues. Also, we focus on acclaimed films⁴ to control for the possible negative influence of bad content (speakers may still dislike the movies, though). YouTube aims to control both content and dialogue.

Finally, we explore the linguistic Corpus of Regional African American Language (CORAAL) [38] as representations of spoken African-American English. For comparisons, we employ the Buckeye [50] dataset focused on Caucasian⁵ speakers from central Ohio. Both datasets are focused on spoken interviews with transcripts. Buckeye was recommended to us by the curators of CORAAL.

We now discuss how we identified African-American English expressions (AAE expressions). As stated, we manually transcribed the Black Talk dictionary [54]. Since AAE first emerged as an oral language, the main intent of this dictionary was not to define the etymological history of terms; instead, it concentrates on the meanings and significance of expressions. The organizers of the not-yet-published Oxford Dictionary of African American English referred us to the Black Talk dictionary as a reliable source.

The Black Talk dictionary comprises more than 1800 entries. Since some entries are sentences instead of single terms, they may apply to different pronouns. In such cases, the possible use cases are listed. For example, "BREAK HIM/HER/THEM OFF SOMETHING" becomes three expressions. Our transcription of the entries considers every possible combination presented.

In Table 2, we present a summary of our datasets in the number of sentences (or utterances), number of words (non-unique), and number of African-American English expressions present. Over the next few subsections, we now detail each dataset.

Twitter: The Twitter dataset comes from the Twitter AAE⁶ website. To create the dataset, the authors [9, 10] developed a Latent Dirichlet Allocation (LDA) based topic model that took into account both the frequency of common terms used in AAE as well as Census data. An initial race estimate is performed based on the location

from which the account was tweeted. This information is combined with the presence of key terms to derive different latent topics for the corpus. Topics were then explored to label AAE and non-AAE.

Although the authors label over 80,000 tweets, we focus on a smaller sample of 500 tweets that the authors manually inspected. These tweets were manually labeled with PoS tags to derive an African-American English PoS model. According to the authors, more than 18% of the terms used within the African-American tweets are not in the standard English dictionary. It is also very common to find words written in their phonological style in AAE - e.g. tha (the), iont (I don't), ova (over), and so on - while the contrary was found to never happen in the Non-AAE tweets.

YouTube: The YouTube dataset is a collection of subtitles from YouTube movie review videos. A single speaker talks to the audience about a movie production listed among the most relevant movies ever. We considered Rotten Tomatoes' top 100 best movies of all time ranking due to their prestige among the audience and because they have a higher probability of being well-spoken in a review. For each of the top 50 movies from the ranking, two authors manually searched and cataloged as many videos as possible. The authors determined Demographic labels, namely, African-American Women, African-American Men, non-African-American Women, and non-African-American Men, in order of appearance when querying the movie name on YouTube. Since YouTube doesn't naturally disclose demographic information about its users, we had to restrict our search only to producers who happened to appear on the screen at least once throughout the entire video. The list of movies and the respective YouTube channels is available⁷.

When publishing videos on YouTube, the creators can either explicitly inform their videos' subtitles or let the YouTube transcription model automatically caption them. Nonetheless, differently from manually informed subtitles, the captioning mechanism, by default⁸. Only 21 videos had manual subtitles in our dataset⁹. Thus, for fair comparisons, we use automatic transcriptions as these are available in **every** video. Finally, it is important to state that transcriptions are not punctuated by YouTube. We use ML models to correct this behavior on the YouTube dataset and the CORAAL/Buckeye dataset (below).

Considering the observational nature of our study, an extensive effort was applied to control the confounding variables' effect on the conclusions. The selection of the most prestigious movies of all time was an attempt to reduce the chance of having negative reviews, which would comprise higher scores in the toxicity analysis. We also tried to find at least one single-speaker video review for every movie to reduce any sampling bias impact. More importantly, the reviews' first-person nature helps eliminate the possibility of other people's opinions influencing the argumentative paths.

On author inferred demographic variables: It is worth noticing that identifying race/gender is subjective and prone to errors - we only have our view and not the content producer's identification. Thus, we avoid using YouTube demographic variables as input in **any** analysis. We employed author-inferred demographics to collect a diverse (to the extent possible) dataset of utterances. Instead,

⁴<https://www.rottentomatoes.com/top/bestofrt/>

⁵<https://buckeyecorpus.osu.edu/php/corpusinfo.php>

⁶<http://slanglab.cs.umass.edu/TwitterAAE/>

⁷https://anonymous.4open.science/r/aae_bias-D396/data/youtube_data_description.csv

⁸<https://support.google.com/youtube/thread/70343381>

⁹Results do not change if we look only into these few videos.

we rely on the # of AAE Expressions as an explanatory variable. Nevertheless, we do see an increase in AAE expressions based on the inferred demographic on YouTube (see Table 2).

CORAAL and Buckeye: The Corpus of Regional African American Language [38] is a long-term corpus developed and maintained by the University of Oregon with the support of the National Science Foundation. The dataset comprises more than 150 socio-linguistic interviews with African-American English speakers born between 1891 and 2005. The dataset contains the orthographic transcriptions of interviews, together with the person’s age, gender, and city they live in. Thus, each interview from the corpus encompasses many subjects from a given city/community.

Unlike the YouTube data, the transcriptions here represent the entire sentence, accounting for complete punctuation, line-level notes, and even non-linguistic sounds. Beyond that, the data also tracks the interviewer’s voice in the dialog. The interviews allow the speakers to talk freely about different topics, an interesting feature that emulates the diversity of daily interactions and mood variations. The dataset aggregates five major sub-corpora from different locations in the United States of America, namely, Atlanta (2017), Washington (1968 and 2016), Lower East Side (2009), Princeville (2004), Rochester (2016), and Valdosta (2017).

The Buckeye [50] corpus is an effort started in 1999 and supported by the National Institute on Deafness and Other Communication Disorders and the Office of Research at Ohio State University. The initial goal was to gather approximately 300,000 words of speech conversation from central Ohio speakers, keeping track of time and phonetic information. To reach that objective, researchers selected a group of 40 middle-class Caucasian speakers.

Similar to the YouTube dataset, Buckeye sentences are not punctuated. However, instead of automatically generated captions, this corpus employed transcribers who were explicitly instructed not to use punctuation within the utterances and not to try to correct possible speech “errors” (we segmented sentences ourselves).

3.0.1 Sentence Segmentation. Except for CORAAL, the datasets don’t necessarily follow the correct orthographic rules about punctuation. Considering the other two transcribed corpus (i.e. YouTube and Buckeye), we should expect their sentences to be segmented not according to their inherent meaning but to silent intervals (not necessarily long ones) after a continuous pronunciation of words. Such segmentation can drastically misrepresent the sentences’ meanings and consequently derive misleading conclusions about the data.

To reduce the impacts of incorrect segmentation in later analysis, we employed a machine learning-based segmentation to all corpus, except the one from Twitter. We believe tweets are self-contained messages where punctuation is not necessarily crucial to the audience’s understanding. Consequently, segmentation is not necessary. On the other hand, we segment the only correctly punctuated corpus, CORAAL. Since we intend to compare the CORAAL dataset directly against Buckeye’s, we should try to reduce the confounding factors (segmentation). The segmentation task was performed using NVIDIA’s NeMo Toolkit¹⁰.

3.0.2 On the Impact of Swear Words. We executed two versions of our experiments, one considering utterances with swear words and

another without. The swear words we considered were taken from the No Swearing project¹¹, a cooperative effort to help programmers remove unwanted language from their applications. At the time of writing, the project listed 363 curse words. Overall, we found no statistical difference in our results nor any significant difference in our figures. Here, we present results with swear words.

3.0.3 Linguistic Features. One of the most relevant aspects of our analysis is directly derived from controlling for linguistic and grammatical features from the available utterances. Thus, our research focuses on word classes, or Part-of-Speech (PoS), (e.g., Verb, Noun, Adjective, etc.) and language dimensions (e.g., Anger, Hate, etc.).

The PoS features consider the function of each word in the sentence. The word *smile* can be considered a verb; however, it can also be considered an adjective when used in certain scenarios, as in “*The smiling baby is really cute*”. This information can help us understand the sentence’s composition regarding word classes. Only on TwitterAAE is it that we have manual PoS features. For such a reason and for fair comparisons, we use automatic PoS features in all datasets. Thus, to classify the tokens according to their PoS categories, we employ a black-box model¹². We also point out that similar results arise when using manual PoS tags on TwitterAAE.

To define linguistic features, we used the Linguistic Inquiry and Word Count (LIWC) software [49] in its 2015 release. LIWC is a research effort that maps words to psychological features (i.e., language dimensions) of speech. A single word may be assigned to as many suitable categories as necessary. For example, the word *cried* is a 10-categories term (e.g. Affect, Positive Tone, Emotion, Negative Emotion, Sad Emotion, Verbs, Past Focus, etc.). Features are computed based on the number of occurrences.

4 RESULTS

We now present our results. Initially, we compare the model scores for utterances with and without AAE expressions. Next, we determine which factors impact the outcome of different methods. Our final analysis compares semantically similar sentences.

4.1 Scores Per Usage of AAE Expressions

Figure 1a compares toxicity/sentiment scores for utterances with and without AAE expressions. The figure shows the complementary cumulative distributions (CCDF) considering the number of AAE expressions on the utterance. That is, the x-axis of the figure shows the score, whereas the y-axis captures the fraction of sentences for that with scores greater than the one on the x-axis. In this and the other CDFs presented in our study, some care must be taken when interpreting results from Textblob and Vader. For these methods, negative values point toward negative sentiment, whereas positive values point toward positive sentiment. For the other models, 0 commonly indicates a not-negative (Vader) sentiment or non-toxic, whereas 1 is a negative sentiment or toxic utterance. This difference comes from sentiment analysis methods commonly measure *polarity* from -1 to 1. Unlike Textblob and Flair, Vader returns the probability of a negative, positive, or neutral score (adding up to one). Here, we focus on the negative probability.

¹⁰<https://github.com/NVIDIA/NeMo>

¹¹<https://www.noswearing.com/>

¹²<https://spacy.io>

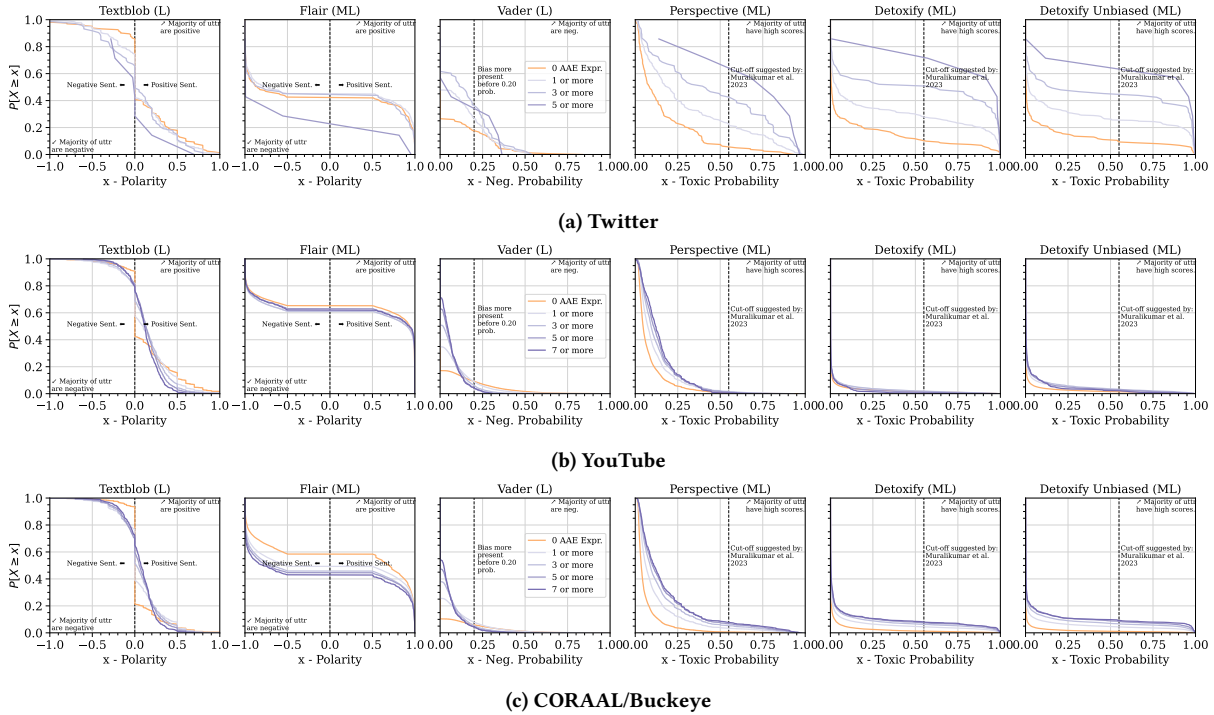


Figure 1: Score distributions for sentences with and without AAE Expressions

We can initially see that utterances with AAE expressions receive much higher scores for Twitter and toxicity models (Perspective, Detoxify, and Detoxify Unbiased). There is a clear tendency for statistical dominance – the CDF, or y-value, of utterances with AAE expressions is above the one without AAE expressions regardless of x-values. This behavior is attenuated as more AAE expressions are considered in the utterance (we considered those with at least one, three, five, or seven expressions). For instance, considering Toxicity models on Twitter, the highest 20% scoring utterances *without* AAE expressions achieve a minimum of 0.25. For utterances *with at least one* AAE expression, this minimum is around 0.55.

Using a Kolmogorov-Smirnov test, we compared whether each CDF with AAE expressions differs from those without. Under $p < 0.01$, Flair did not show this difference on the Twitter dataset when considering utterances with at least one expression. This is the **only** case where we failed to reject the null hypothesis. When we consider lexical models, Textblob and Vader, it appears they present a mild bias on lower-scoring sentences. To help understand this issue, focus on the sentences with scores below 0.5 for Textblob and 0.20 for Vader. Whether this is an issue will depend on the cut-off developers employ (e.g., ± 0.5 appears to mitigate the issue).

From Figures 1b and 1c, we can see the same trend that occurred on Twitter for the lexical approaches (Textblob and Vader), also occurs on YouTube and CORAAL/Buckeye. This finding likely stems from the fact that such approaches employ manually curated rules that do not consider AAE, a positive aspect of these approaches. On these datasets, Flair is biased, achieving lower scores (polarity or tendency to rate as more negative) for utterances with AAE expressions. The ML models still present biases on YouTube (Figure 1b).

However, it is important to point out two facts: (1) albeit statistically significant, this bias is negligible on Detoxify for YouTube; (2) such bias is less present when considering the suggested cut-off of 0.55 [45] for toxicity models. Biases **may not** be an issue when using large cut-off values. To further investigate biases, we now correlate scores with linguistic features.

4.2 Impact of Grammatical/Linguistic Features

In Table 3, we present our Logistic regression results for Twitter. For each tweet, we counted the number of AAE expressions, LIWC categories, and PoS tags as features. Being counts, all of these features have positive values only. To present regression coefficients on a similar, features were *Min-Max* scaled to the [0, 1] range before the regression was executed. Models were executed with an intercept variable and no regularization. Models that output polarity had such polarity values re-scaled to [0, 1] also (by adding one and dividing by two). The table presents only the statistically significant features ($p < 0.05$). LIWC features are identified by LIWC_, and PoS tags by POS_. The demographic variable (AAE tweet or non-AAE tweet) was used as a Twitter and CORAAL/Buckeye feature.

Similar results are found for the other two datasets. For simplicity, instead of presenting the tables with such results, we shall discuss these results throughout the text. We also note similar results when using the manual POS tags for TwitterAAE (omitted for space).

Our feature of most interest is the AAE_EXPR (the number of AAE expressions on the utterance). The other features act as control variables to ensure that, on some level, such expressions are not being confounded with other grammatical/linguistic attributes.

Features	Textblob (L)	Flair (ML)	Vader (L)	Perspective (ML)	Detoxify (ML)	Detoxify U. (ML)
	👎 < 0	👍 > 0			👎 > 0	👍 < 0
AAE_EXPR			0.0934	0.2238	0.1779	
LIWC_SWEAR				0.8567	0.9492	1.2792
LIWC_SEXUAL	-0.3354	-1.3418		0.4657	0.5942	0.5609
LIWC_NETSPEAK		-2.6379		-0.4496	-0.8121	-0.9239
LIWC_INFORMAL		2.5228		0.4386	0.7988	0.946
LIWC_NEGATE		-0.8331		0.2075		0.1855
LIWC_FILLER		-2.1683			-0.6053	-0.6377
LIWC_ASSENT				-0.1599	-0.2614	-0.1985
LIWC_MALE	-0.3237			-0.1505	-0.1968	-0.191
POS_X	0.5655		-0.1636	-0.3857	-0.4255	-0.3066
POS_DET				0.1623	0.3369	0.3268
DEMOGRAPHIC				0.0508	0.1262	0.0762

Table 3: Logistic Regression coefficients for the Twitter dataset with $p < 0.05$. Each model’s five most relevant coefficients are presented in bold, whereas not statistically significant coefficients were omitted. When a coefficient pushes towards a negative sentiment or toxic score, we color it red (👎). Positive sentiment and non-toxic score is colored green (👍).

From the table, we can see that this feature pushes the polarity of the Vader model towards having fewer sentiments. The feature is not significant for the other sentiment analysis models. Nevertheless, this feature is significant for Perspective and Detoxify (but not for Detoxify Unbiased). When we consider YouTube, AAE_EXPR is statistically significant for Vader (0.0617), Perspective (0.2488), and –surprisingly– Detoxify Unbiased (0.1334). On CORAL/Buckeye, it was significant for Flair (-0.9432), Perspective (0.3291), Detoxify (0.1754), and Detoxify Unbiased (0.2126).

For Twitter and CORAAL/Buckeye, our demographic variable (DEMOGRAPHIC) was used as a categorical feature. This feature was statistically significant for Twitter and not CORAAL/Buckeye. Twitter is the only dataset where the demographic variable was developed to align with AAE utterances. On CORAAL/Buckeye, an African-American may not employ AAE, or a Caucasian may employ AAE (in fact, the usage of expressions is comparable in Table 2).

Considering the other features, some linguistic features are expected to push models towards negative sentiment or toxic scores (this is the case for the feature LIWC_SWEAR in every dataset). Finally, PoS features were only statistically significant for Twitter: quantifiers, POS_DET, and the unknown/other tag, POS_X.

4.3 Semantic Comparison

We now seek to answer the following question: *How do highly semantically similar pairs of utterances that achieve diverging scores differ in their usage of AAE expressions?* Notice that we have two conditions here: (1) being similar in meaning and (2) achieving diverging scores. Such pairs of utterances are interesting because they control for confounding factors in semantics.

We employ a large language model as our semantic feature extractor [62]. With this model, utterances are mapped to an embedding vector. We compare these vectors using a cosine similarity for pairs of utterances: -1 indicates entirely dissimilar, 0 indicates a lack of relationship, and 1 shows completely similar. We deem two sentences similar when the cosine score is above 0.5 (this is less than 0.1% of pairs as discussed below).

For each dataset and sentiment/toxicity method of our study, we isolated the top 2.5% and bottom 2.5% scoring utterances. To perform a single analysis, we standardized scores. For Perspective

and Detoxify, the top 2.5% have a higher chance of being toxic, whereas for Vader, Flair, and Textblob, the top 2.5% have negative polarity. Due to memory constraints sampled 100,000 of such pairs per dataset and method. Next, we focused only on pairs where one of the sentences had at least one AAE expression. If this is not the case, the difference in score certainly is not due to the number of AAE expressions employed. When this is the case, the usage of AAE expressions may be the underlying cause. Due to the small sample size, we did not find any pairs on Twitter that met our conditions.

Combining every method, we found 585,679 unique pairs on YouTube with diverging scores (our top 2.5% versus the bottom 2.5%). Out of these, 568 had a cosine similarity above 0.5%. On CORAAL/Buckeye, we found 592,606 unique pairs from our diverging scores filter, with 243 being highly similar. For each setting (YouTube or CORAAL/Buckeye), we computed the number of pairs where (d_{bias}) the most toxic or most negative had more AAE expressions; (d_{eq}) both had the same amount of expressions, and (d_{no_bias}) the least toxic or most positive had more expressions.

For YouTube, we have that: $d_{bias} = 272$ (48%), $d_{eq} = 87$ (15%), and $d_{no_bias} = 209$ (37%). Whereas on CORAAL/Buckeye: $d_{bias} = 187$ (77%), $d_{eq} = 26$ (11%), and $d_{no_bias} = 30$ (12%). Under a Binomial test and $p < 0.01$, in both cases, we find that $d_{bias} > d_{no_bias}$. Thus, results show statistical evidence of bias¹³.

5 CONCLUSIONS AND LIMITATIONS

This paper investigates the biases of sentiment analysis/toxicity methods regarding the usage of AAE expressions. We analyzed the performance of six well-known off-the-shelf methods in light of four different datasets. Our datasets ranged from online texts from Twitter to single-speaker closed captions from YouTube and spoken English encompassing daily live situations.

Considering the latter, or nonexistent, introduction of AAE in ML datasets, the under-representation of such expressions leads ML models to present a systemic bias towards AAE. We argue that the biggest problems derive directly from the absence of context in the utterances. Since they employ human-crafted rules, lexical-based (rule) approaches tend to be less biased than ML models.

¹³ d_{eq} is not considered as the counts of AAE expressions usage are certainly not the issue in these settings (they are equal)

ETHICAL CONSIDERATIONS

Ethical Concerns: One of the ethical concerns of our study comes from using demographic variables related to race **without** ground-truth self-identification labels for speakers. To mitigate this issue, we refrained from using author-inferred demographic variables in our study (on the YouTube dataset). CORAAL/Buckeye are well-established in linguistics, with CORAAL focusing solely on African Americans. This issue is not present on Twitter, as labels come from using AAE or not (regardless of race). We also point out that our main statistical variable of study is not race. We focus on the usage of AAE expressions, where such expressions came from reliable and suggested sources (by the organizers of a well-known dictionary).

Unintended Impact: Readers may interpret our research as against ML models or automatic utterance scoring tools. We point out that this is **not** our statement. Our research advances both recent and large literature on the unintended biases of Lexical/AI/ML models. We hope our findings will improve how such tools are used; model advances towards fewer biases or both.

Researcher Background: The majority of authors of this study are from a region where racial discrimination is still very present in the population's day-to-day lives. Our research aims to foster the ongoing discussion on how AI impacts the lives of different historically segregated communities.

If readers deem any terms or expressions used in this paper offensive, we point out that it was not deliberate.

REFERENCES

- [1] Abubakar Abid, Maheen Farooqi, and James Zou. 2021. Large language models associate Muslims with violence. *Nature Machine Intelligence* 3, 6 (2021), 461–463.
- [2] CJ Adams. 2018. New York Times: Using AI to host better conversations. <https://blog.google/technology/ai/new-york-times-using-ai-host-better-conversations/>.
- [3] Alan Akbik, Tanja Bergmann, Duncan Blythe, Kashif Rasul, Stefan Schweter, and Roland Vollgraf. 2019. FLAIR: An easy-to-use framework for state-of-the-art NLP. In *NAACL 2019, 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*. 54–59.
- [4] Marzieh Babaeianjelodar, Stephen Lorenz, Josh Gordon, Jeanna Matthews, and Evan Freitag. 2020. Quantifying gender bias in different corpora. In *Companion Proceedings of the Web Conference 2020*. 752–759.
- [5] Arne F. Ball. 1992. Cultural preference and the expository writing of African-American adolescents. *Written Communication* 9, 4 (1992), 501–532.
- [6] Ari Ball-Burack, Michelle Seng Ah Lee, Jennifer Cobbe, and Jatinder Singh. 2021. Differential tweetment: Mitigating racial dialect bias in harmful tweet detection. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 116–128.
- [7] David Bamman, Chris Dyer, and Noah A Smith. 2014. Distributed representations of geographically situated language. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 828–834.
- [8] John Baugh. 1981. Runnin' down Some Lines: The Language and Culture of Black Teenagers.
- [9] Su Lin Blodgett, Lisa Green, and Brendan O'Connor. 2016. Demographic dialectal variation in social media: A case study of African-American English. *arXiv preprint arXiv:1608.08868* (2016).
- [10] Su Lin Blodgett, Johnny Wei, and Brendan O'Connor. 2018. Twitter universal dependency parsing for African-American and mainstream American English. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1415–1425.
- [11] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems* 29 (2016).
- [12] Zhenpeng Chen, Jie M Zhang, Max Hort, Federica Sarro, and Mark Harman. 2022. Fairness testing: A comprehensive survey and analysis of trends. *arXiv preprint arXiv:2207.10223* (2022).
- [13] Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the international AAAI conference on web and social media*, Vol. 11. 512–515.
- [14] Mark Diaz, Isaac Johnson, Amanda Lazar, Anne Marie Piper, and Darren Gergle. 2018. Addressing age-related bias in sentiment analysis. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–14.
- [15] Joey Lee Dillard. 1977. *Lexicon of Black English*. ERIC.
- [16] Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. 67–73.
- [17] Jacob Eisenstein, Brendan O'Connor, Noah A Smith, and Eric P Xing. 2014. Diffusion of lexical change in social media. *PLoS one* 9, 11 (2014), e113114.
- [18] Anjalie Field, Su Lin Blodgett, Zeerak Waseem, and Yulia Tsvetkov. 2021. A Survey of Race, Racism, and Anti-Racism in NLP. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 1905–1925.
- [19] Sarah Florini. 2014. Tweets, Tweeps, and Signifyin' Communication and Cultural Performance on "Black Twitter". *Television & New Media* 15, 3 (2014), 223–237.
- [20] Patricia Friedrich. 2020. When Englishes go digital. *World Englishes* 39, 1 (2020), 67–78.
- [21] Patricia Friedrich and Eduardo Diniz de Figueiredo. 2016. *The sociolinguistics of digital Englishes*. Routledge.
- [22] David Garcia, Ingmar Weber, and Venkata Rama Kiran Garimella. 2014. Gender asymmetries in reality and fiction: The bechdel test of social media. In *Eighth International AAAI Conference on Weblogs and Social Media*.
- [23] Anastasia Giachanou and Fabio Crestani. 2016. Like it or not: A survey of twitter sentiment analysis methods. *ACM Computing Surveys (CSUR)* 49, 2 (2016), 1–41.
- [24] Alec Go, Richa Bhayani, and Lei Huang. 2009. Twitter sentiment classification using distant supervision. *CS224N project report, Stanford* 1, 12 (2009), 2009.
- [25] A Gomes, D Antoniali, and T Dias-Oliva. 2019. Drag queens and artificial intelligence. Should computers decide what is toxic on the internet. *Internet Lab blog* (2019).
- [26] Hila Gonen and Yoav Goldberg. 2019. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. *arXiv preprint arXiv:1903.03862* (2019).
- [27] Mark Graham, Bernie Hogan, Ralph K Straumann, and Ahmed Medhat. 2014. Uneven geographies of user-generated information: Patterns of increasing informational poverty. *Annals of the Association of American Geographers* 104, 4 (2014), 746–764.
- [28] Lisa J Green. 2002. *African American English: a linguistic introduction*. Cambridge University Press.
- [29] Tommi Gröndahl, Luca Pajola, Mika Juuti, Mauro Conti, and N Asokan. 2018. All you need is "love" evading hate speech detection. In *Proceedings of the 11th ACM workshop on artificial intelligence and security*. 2–12.
- [30] Laura Hanu and Unitary team. 2020. Detoxify. Github. <https://github.com/unitaryai/detoxify>.
- [31] Camille Harris, Matan Halevy, Ayanna Howard, Amy Bruckman, and Diyi Yang. 2022. Exploring the role of grammar and word choice in bias toward african american english (aae) in hate speech classification. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. 789–798.
- [32] Hossein Hosseini, Sreeram Kannan, Baosen Zhang, and Radha Poovendran. 2017. Deceiving google's perspective api built for detecting toxic comments. *arXiv preprint arXiv:1702.08138* (2017).
- [33] Mingqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. 168–177.
- [34] Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. 2020. Unintended machine learning biases as social barriers for persons with disabilities. *ACM SIGACCESS Accessibility and Computing* (2020), 1–1.
- [35] Clayton Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, Vol. 8. 216–225.
- [36] Sen Jia, Thomas Lansdall-Welfare, and Nello Cristianini. 2015. Measuring gender bias in news images. In *Proceedings of the 24th International Conference on World Wide Web*. 893–898.
- [37] Jigsaw. [n. d.]. Perspective API. <https://perspectiveapi.com/>. Accessed: 2023-01-30.
- [38] Tyler Kendall and Charlie Farrington. 2021. The Corpus of Regional African American Language (Version 2021.07). Eugene, OR: The Online Resources for African American Language Project.
- [39] Svetlana Kiritchenko and Saif M Mohammad. 2018. Examining Gender and Race Bias in Two Hundred Sentiment Analysis Systems. *NAACL HLT 2018* (2018), 43.
- [40] Animesh Koratana and Kevin Hu. 2018. Toxic speech detection. URL: <https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1194> (2018).
- [41] Deepak Kumar, Patrick Gage Kelley, Sunny Consolvo, Joshua Mason, Elie Bursztein, Zakir Durumeric, Kurt Thomas, and Michael Bailey. 2021. Designing Toxic Content Classification for a Diversity of Perspectives.. In *SOUPS@ USENIX Security Symposium*. 299–318.

- [42] Steven Loria. 2018. textblob Documentation. *Release 0.15 2* (2018).
- [43] Patricia Georgiou Marie Pellat. 2018. Perspective Launches In Spanish With El País. <https://medium.com/jigsaw/perspective-launches-in-spanish-with-el-pa%C3%ADs-dc2385d734b2>.
- [44] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [45] Meena Devii Muralikumar, Yun Shan Yang, and David W McDonald. 2023. A Human-Centered Evaluation of a Toxicity Detection API: Testing Transferability and Unpacking Latent Attributes. *ACM Transactions on Social Computing* (2023).
- [46] Lisa Nakamura. 2013. *Cybertypes: Race, ethnicity, and identity on the Internet*. Routledge.
- [47] Nikolaos Pappas, Georgios Katsimpras, and Efstathios Stamatatos. 2013. Distinguishing the popularity between topics: a system for up-to-date opinion retrieval and mining in the web. In *Computational Linguistics and Intelligent Text Processing: 14th International Conference, CICLing 2013, Samos, Greece, March 24-30, 2013, Proceedings, Part II 14*. Springer, 197–209.
- [48] Daniel Borkan Patricia Georgiou, Marie Pellat. 2019. Parlons-en! Perspective and Tune are now available in French. <https://medium.com/jigsaw/perspective-tune-are-now-available-in-french-c4cf1ca198f2>.
- [49] 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), 2001.
- [50] Mark A Pitt, Keith Johnson, Elizabeth Hume, Scott Kiesling, and William Raymond. 2005. The Buckeye corpus of conversational speech: Labeling conventions and a test of transcriber reliability. *Speech Communication* 45, 1 (2005), 89–95.
- [51] Filipe N Ribeiro, Matheus Araújo, Pollyanna Gonçalves, Marcos André Gonçalves, and Fabricio Benevenuto. 2016. Sentibench—a benchmark comparison of state-of-the-practice sentiment analysis methods. *EPJ Data Science* 5, 1 (2016), 1–29.
- [52] Max Roser, Hannah Ritchie, and Esteban Ortiz-Ospina. 2015. Internet. *Our World in Data* (2015). <https://ourworldindata.org/internet>.
- [53] Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A Smith. 2019. The risk of racial bias in hate speech detection. In *Proceedings of the 57th annual meeting of the association for computational linguistics*. 1668–1678.
- [54] Geneva Smitherman. 2000. *Black talk: Words and phrases from the hood to the amen corner*. Houghton Mifflin Harcourt.
- [55] Kaikai Song, Ting Yao, Qiang Ling, and Tao Mei. 2018. Boosting image sentiment analysis with visual attention. *Neurocomputing* 312 (2018), 218–228.
- [56] Ezekiel Soremekun, Sakshi Udeshi, and Sudipta Chattopadhyay. 2022. Astraea: Grammar-based fairness testing. *IEEE Transactions on Software Engineering* 48, 12 (2022), 5188–5211.
- [57] Maite Taboada, Julian Brooke, Milan Tofiloski, Kimberly Voll, and Manfred Stede. 2011. Lexicon-based methods for sentiment analysis. *Computational linguistics* 37, 2 (2011), 267–307.
- [58] Rachael Tatman. 2017. Gender and dialect bias in YouTube’s automatic captions. In *Proceedings of the first ACL workshop on ethics in natural language processing*. 53–59.
- [59] Mike Thelwall. 2014. Heart and soul: Sentiment strength detection in the social web with sentistrength, 2017. *Cyberemotions: Collective emotions in cyberspace* (2014).
- [60] Pranav Narayanan Venkit and Shomir Wilson. 2021. Identification of bias against people with disabilities in sentiment analysis and toxicity detection models. *arXiv preprint arXiv:2111.13259* (2021).
- [61] Hao Wang, Doğan Can, Abe Kazemzadeh, François Bar, and Shrikanth Narayanan. 2012. A system for real-time twitter sentiment analysis of 2012 us presidential election cycle. In *Proceedings of the ACL 2012 system demonstrations*. 115–120.
- [62] Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems* 33 (2020), 5776–5788.
- [63] Maciej Widański. 2015. *African American slang: A linguistic description*. Cambridge University Press.
- [64] Theresa Wilson, Paul Hoffmann, Swapna Somasundaran, Jason Kessler, Janyce Wiebe, Yejin Choi, Claire Cardie, Ellen Riloff, and Siddharth Patwardhan. 2005. OpinionFinder: A system for subjectivity analysis. In *Proceedings of HLT/EMNLP 2005 Interactive Demonstrations*. 34–35.
- [65] Austin P Wright, Omar Shaikh, Haekyu Park, Will Epperson, Muhammed Ahmed, Stephane Pinel, Duen Horng Chau, and Diyi Yang. 2021. RECAST: Enabling user recourse and interpretability of toxicity detection models with interactive visualization. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–26.
- [66] Ali Yadollahi, Ameneh Gholipour Shahraki, and Osmar R Zaiane. 2017. Current state of text sentiment analysis from opinion to emotion mining. *ACM Computing Surveys (CSUR)* 50, 2 (2017), 1–33.
- [67] Min Yang, Qiang Qu, Xiaojun Chen, Chaoxue Guo, Ying Shen, and Kai Lei. 2018. Feature-enhanced attention network for target-dependent sentiment classification. *Neurocomputing* 307 (2018), 91–97.