

Towards a Deep Multi-layered Dialectal Language Analysis: A Case Study of African-American English

Jamell Dacon

Michigan State University
East Lansing, MI, USA
daconjam@msu.edu

Abstract

Currently, natural language processing (NLP) models proliferate language discrimination leading to potentially harmful societal impacts as a result of biased outcomes. For example, part-of-speech taggers trained on Mainstream American English (MAE) produce non-interpretable results when applied to African American English (AAE) as a result of language features not seen during training. In this work, we incorporate a human-in-the-loop paradigm to gain a better understanding of AAE speakers' behavior and their language use, and highlight the need for dialectal language inclusivity so that native AAE speakers can extensively interact with NLP systems while reducing feelings of disenfranchisement.

1 Introduction

Over the years, social media users have leveraged online conversational platforms to perpetually express themselves online. For example, African American English (AAE)¹, an English language variety is often heavily used on Twitter (Field et al., 2021; Blodgett et al., 2020). This dialect continuum is neither spoken by *all* African Americans or individuals who identify as BIPOC (Black, Indigenous, or People of Color), nor is it spoken *only* by African Americans or BIPOC individuals (Field et al., 2021; Bland-Stewart, 2005). In some cases, AAE, a low-resource language (LRL) may be the first (or dominant) language, rather than the second (or non-dominant) language of an English speaker.

Specifically, AAE is a regional dialect continuum that consists of a distinct set of lexical

¹A dialectal continuum previously known as Northern Negro English, Black English Vernacular (BEV), Black English, African American Vernacular English (AAVE), African American Language (AAL), Ebonics, and Non-standard English (Labov, 1975; Bailey et al., 1998; Green, 2002, 2014; Baugh, 2008; Bland-Stewart, 2005; King, 2020). It is often referred to as African American Language (AAL) and African American English (AAE). In this work, we use the denotation AAE.

items, some of which have distinct semantic meanings, and may possess different syntactic structures/patterns than in Mainstream American English (MAE) (*e.g.*, differentiating habitual *be* and non-habitual *be* usage) (Stewart, 2014; Dorn, 2019; Jones, 2015; Field et al., 2021; Bland-Stewart, 2005; Baugh, 2008; Blodgett et al., 2020; Labov, 1975). In particular, Green (2002) states that AAE possesses a morphologically invariant form of the verb that distinguishes between habitual action and currently occurring action, namely *habitual be*. For example, “the habitual *be*” experiment² by University of Massachusetts Amherst’s Janice Jackson.

However, AAE is perceived to be “bad english” despite numerous studies by socio/raciolinguists and dialectologists in their attempts to quantify AAE as a legitimized language (Baugh, 2008; Field et al., 2021; Bland-Stewart, 2005; Labov, 1975).

*“[T]he common misconception [is] that language use has primarily to do with words and what they mean. It doesn’t. It has primarily to do with people and what **they** mean.”* – Clark and Schober (1992)

Recently, online AAE has influenced the generation of resources for AAE-like text for natural language (NLP) and corpus linguistic tasks *e.g.*, part-of-speech (POS) tagging (Jørgensen et al., 2016; Blodgett et al., 2018), language generation (Groenwold et al., 2020) and automatic speech recognition (Dorn, 2019; Tatman and Kasten, 2017). POS tagging is a token-level text classification task where each token is assigned a corresponding word category label (see Table 1). It is an enabling tool for NLP applications such as a syntactic parsing, named entity recognition, corpus linguistics, etc. In this work, we incorporate a human-in-the-loop paradigm by directly involving affected (user) communities to understand context and word ambigu-

²<https://www.umass.edu/synergy/fall198/ebonics3.html>

MAE	Input	I have never done this before
	Output	(I, <PRP>), (have, <VBP>), (never, <RB>), (done, <VBN>), (that, <IN>), (before, <IN>)
AAE	Input	I aint neva did dat befo
	Output	(I, <PRP>), (aint, <VBP>), (neva, <NN>), (did, <VBD>)(dat, <JJ>), (befo, <NN>)

Table 1: An illustrative example of POS tagging of semantically equivalent sentences written in MAE and AAE. Each blue and red highlight corresponds to linguistics features of AAE lexical items, and their misclassified NLTK (inferred) tags, respectively.

ities in an attempt to study dialectal language inclusivity in NLP language technologies that are generally designed for dominant language varieties. Dacon and Liu (2021) state that,

“NLP systems aim to [learn] from natural language data, and mitigating social biases become a compelling matter not only for machine learning (ML) but for social justice as well.”

To address these issues, we aim to empirically study *predictive bias* (see Swinton (1981) for definition) *i.e.*, if POS tagger models make predictions dependent on demographic language features, and attempt a dynamic approach in data-collection of non-standard spellings and lexical items. To examine the behaviors of AAE speakers and their language use, we first collect variable (morphological and phonological) rules of AAE language features from literature (Labov, 1975; Bailey et al., 1998; Green, 2002; Bland-Stewart, 2005; Stewart, 2014; Blodgett et al., 2016; Elazar and Goldberg, 2018; Baugh, 2008; Green, 2014) (see Appendix C). Then, we employ 5 trained sociolinguist Amazon Mechanical Turk (AMT) annotators³ who identify as bi-dialectal dominant AAE speakers to address the issue of lexical, semantic and syntactic ambiguity of tweets (see Appendix B for annotation guidelines). Next, we incorporate a human-in-the-loop paradigm by recruiting 20 crowd-sourced diglossic annotators to evaluate AAE language variety (see Table 2). Finally, we conclude by expanding on the need for dialectal language inclusivity.

2 Related Work

Previous works regarding AAE linguistic features have analyzed tasks such as unsupervised domain adaptation for AAE-like language (Jørgensen et al., 2016), detecting AAE syntax (Stewart, 2014), language identification (Blodgett and O’Connor, 2017), voice recognition and transcription (Dorn,

³A HIT approval rate $\geq 95\%$ was used to select 5 bi-dialectal AMT annotators between the ages of 18 - 55, and completed $> 10,000$ HITs and located within the United States.

2019), dependency parsing (Blodgett et al., 2018), dialogue systems (Liu et al., 2020), hate speech/toxic language detection and examining racial bias (Sap et al., 2019; Halevy et al., 2021; Xia et al., 2020; Davidson and Bhattacharya, 2020; Zhou et al., 2021; Mozafari et al., 2020; Xu et al., 2021; Koenecke et al., 2020), and language generation (Groenwold et al., 2020). These central works are conclusive for highlighting systematic biases of natural language processing (NLP) systems when employing AAE in common downstream tasks.

Although we mention popular works incorporating AAE, this dialectal continuum has been largely ignored and underrepresented by the NLP community in comparison to MAE. Such lack of language diversity cases constitutes technological inequality to minority groups, for example, by African Americans or BIPOC individuals, and may intensify feelings of disenfranchisement due to monolingualism. We refer to this pitfall as the *inconvenient truth i.e.*,

“[I]f the systems show discriminatory behaviors in the interactions, the user experience will be adversely affected.” — Liu et al. (2020)

Therefore, we define fairness as the model’s ability to correctly predict each tag while performing zero-shot transfer via dialectal language inclusivity.

Moreover, these aforementioned works do not discuss nor reflect on the *“role of the speech and language technologies in sustaining language use”* (Labov, 1975; Bird, 2020; Blodgett et al., 2020) as,

“... models are expected to make predictions with the semantic information rather than with the demographic group identity information” — Zhang et al. (2020).

Interactions with everyday items is increasingly mediated through language, yet systems have limited ability to process less-represented dialects such as AAE. For example, a common AAE phrase, *“I had a long ass day”* would receive a lower sentiment polarity score because of the word *“ass”*, a (noun) term typically classified as offensive; however, in AAE, this term is often used as an emphatic, cumulative adjective and perceived as non-offensive.




Figure 1: An illustration of inferred and manually-annotated AAE tag counts from k randomly sampled tweets.

Motivation: We want to test our hypothesis that training each model on correctly tagged AAE language features will improve the model’s performance, interpretability, explainability, and usability to reduce predictive bias.

3 Dataset and Annotation

3.1 Dataset

We collect 3000 demographically-aligned African American (AA) tweets possessing an average of 7 words per tweet from the publicly available TwitterAAE corpus by Blodgett et al. (2016). Each tweet is accompanied by inferred geolocation topic model probabilities from Twitter + Census demographics and word likelihoods to calculate demographic dialect proportions. We aim to minimize (linguistic) discrimination by sampling tweets that possess over 99% confidence to develop “fair” NLP tools that are originally designed for dominant language varieties by integrating non-standardized varieties. More information about the TwitterAAE dataset, including its statistical information, annotation process, and the link(s) to downloadable versions can be found in Appendix A.

3.2 Preprocessing

As it is common for most words on social media to be plausibly semantically equivalent, we denoise each tweet as tweets typically possess unusual spelling patterns, repeated letter, emoticons and emojis⁴. We replace sequences of multiple

⁴Emoticons are particular textual features made of punctuation such as exclamation marks, letters, and/or numbers

repeated letters with three repeated letters (e.g., *Hmmmmmmmm* → *Hmmm*), and remove all punctuation, “@” handles of users and emojis. Essentially, we aim to denoise each tweet only to capture non-standard spellings and lexical items more efficiently.

3.3 Annotation

First, we employ off-the-shelf taggers such as spacy⁵ and TwitterNLP⁶; however, the Natural Language Toolkit (NLTK) (Loper and Bird, 2002) provides a more fine-grained Penn Treebank Tagset (PTB)⁷ along with evaluation metrics per tag such as F1 score. Next, we focus on aggregating the appropriate tags by collecting and manually-annotating tags from AAE/slang-specific dictionaries to assist the AMT annotators, and later we contrast these aggregated tags with inferred NLTK PTB inferred tags. In Figure 1, we display NLTK inferred and manually-annotated AAE tags from $k = 300$ randomly sampled tweets.

- **The Online Slang Dictionary** (American, English, and Urban slang)⁸ - created in 1996, this is the oldest web dictionary of slang words, neologisms, idioms, aphorisms, jargon, informal speech, and figurative usages. This dic-

to create pictorial icons to display an emotion or sentiment (e.g., “;)” ⇒ *winking smile*), while emojis are small text-like pictographs of faces, objects, symbols, etc.

⁵<https://spacy.io>

⁶<https://github.com/ianozsvald/ark-tweet-nlp-python>

⁷<https://www.guru99.com/pos-tagging-chunking-nltk.html>

⁸<http://onlineslangdictionary.com>

Tags	Category	AAE Example(s)	MAE Equivalent(s)
CC	Coordinating Conjunction	<i>doe/tho, n, bt</i>	<i>though, and, but</i>
DT	Determiner	<i>da, dis, dat</i>	<i>the, this, that</i>
EX	Existential There	<i>dea</i>	<i>there</i>
IN	Preposition/ Conjunction	<i>fa, cuz/cause, den</i>	<i>for, because, than</i>
JJ	Adjective	<i>foine, hawt</i>	<i>fine, hot</i>
PRP	Pronoun	<i>u, dey, dem</i>	<i>you, they, them</i>
PRP\$	Personal Pronoun	<i>ha</i>	<i>her</i>
RB	Adverb	<i>tryna, finna, jus</i>	<i>trying to, fixing to, just</i>
RBR	Adverb, comparative	<i>mo, betta, hotta</i>	<i>more, better, hotter</i>
RP	Particle	<i>bout, thru</i>	<i>about, through</i>
TO	Infinite marker	<i>ta</i>	<i>to</i>
UH	Interjection	<i>wassup, ion, ian</i>	<i>what's up, I don't</i>
VBG	Verb, gerund	<i>sleepin, gettin</i>	<i>sleeping, getting</i>
VBZ	Verb, 3rd-person present tense	<i>iz</i>	<i>is</i>
WDT	Wh-determiner	<i>dat, wat, wus, wen</i>	<i>that, what, what's, when</i>
WRB	Wh-adverb	<i>hw</i>	<i>how</i>

Table 2: Accurately tagged (observed) AAE and English phonological and morphological **linguistic** feature(s) accompanied by their respective MAE equivalent(s).

tionary possesses more than 24,000 real definitions and tags for over 17,000 slang words and phrases, 600 categories of meaning, word use mapping and aids in addressing lexical ambiguity.

- **Word Type**⁹ - an open source POS focused dictionary of words based on the Wiktionary¹⁰ project by Wikimedia¹¹. Researchers have parsed Wiktionary and other sources, including real definitions and categorical POS word use cases necessary to address the issue of lexical, semantic and syntactic ambiguity.

3.4 Human Evaluation

After an initial training of the AMT annotators, we task each annotator to annotate each tweet with the appropriate POS tags. Then, as a calibration study we attempt to measure the inter-annotator agreement (IAA) using Krippendorff’s α . By using NLTK’s (Loper and Bird, 2002) `nltk.metrics.agreement`, we calculate a Krippendorff’s α of 0.88. We did not observe notable distinctions in annotator agreement across the individual tweets. We later randomly sampled 300 annotated tweets and recruit 20 crowd-sourced annotators to evaluate AAE language variety. To recruit 20 diglossic annotators¹², we created a volunteer questionnaire with annotation guidelines, and

⁹<https://wordtype.org/>

¹⁰<https://www.wiktionary.org>

¹¹<https://www.wikimedia.org>

¹²Note that we did not collect certain demographic information such as gender or race, only basic demographics such as age (18-55 years), state and country of residence.

released it on LinkedIn. The full annotation guidelines can be found in Appendix B. Each recruited annotator is tasked to judge sampled tweets and list their MAE equivalents to examine contextual differences of simple, deterministic morphosyntactic substitutions of dialect-specific vocabulary in standard English or MAE texts—a *reverse* study to highlight several varieties of AAE (see Table 2).

4 Methodology

In this section, we describe our approach to perform a preliminary study to validate the existence of predictive bias (Elazar and Goldberg, 2018; Shah et al., 2020) in POS models. We first introduce the POS tagging, and then propose two ML sequence models.

4.1 Part-of-Speech (POS) Tagging

We consider POS tagging as it represents word syntactic categories and serves as a pre-annotation tool for numerous downstream tasks, especially for non-standardized English language varieties such as AAE (Zampieri et al., 2020). Common tags include prepositions, adjective, pronoun, noun, adverb, verb, interjection, etc., where multiple POS tags can be assigned to particular words due to syntactic structural patterns. This can also lead to misclassification of non-standardized words that do not exist in popular pre-trained NLP models.

4.2 Models

We propose to implement two well known sequence modeling algorithms, namely a Bidirectional

Long Short Term Memory (Bi-LSTM) network, a deep neural network (DNN) (Hochreiter and Schmidhuber, 1997; Graves and Schmidhuber, 2005) that has been used for POS tagging (Ling et al., 2015; Plank et al., 2016), and a Conditional Random Field (CRF) (Lafferty et al.) typically used to identify entities or patterns in texts by exploiting previously learned word data.

Taggers: First, we use NLTK (Loper and Bird, 2002) for automatic tagging; then, we pre-define a feature function for our CRF model where we optimized its L1 and L2 regularization parameters to 0.25 and 0.3, respectively. Later, we train our Bi-LSTM network for 40 epochs with an Adam optimizer, and a learning rate of 0.001. Note that each model would be accompanied by error analysis for a 70-30 split of the data with 5-fold cross-validation to obtain model classification reports, for metrics such as precision, recall and F1-score.

5 Operationalization of AAE as an English Language Variety

As (online) AAE can incorporate non-standardized spellings and lexical items, there is an active need for a human-in-the-loop paradigm as humans provide various forms of feedback in different stages of workflow. This can significantly improve the model’s performance, interpretability, explainability, and usability. Therefore, crowd-sourcing to develop language technologies that consider who created the data will lead to the inclusion of diverse training data, and thus, decrease feelings of marginalization. For example, CORAAL¹³, is an online resource that features AAL text data, recorded speech data, etc., into new and existing NLP technologies, AAE speakers can extensively interact with current NLP language technologies.

Consequently, to quantitatively and qualitatively ensure fairness in NLP tools, artificial intelligence (AI) and NLP researchers need to go beyond evaluation measures, word definitions and word order to assess AAE on a token-level to better understand context, culture and word ambiguities. We encourage both AI and NLP practitioners to prioritize collecting a set of relevant labeled training data with several examples of informal phrases, expressions, idioms, and regional-specific varieties. Specifically, in models intended for broad use such as sentiment analysis by partnering with low-resource and di-

alectal communities to develop impactful speech and language technologies for dialect continua such as AAE to minimize further stigmatization of an already stigmatized minority group.

6 Conclusion

Throughout this work, we highlight the need to develop language technologies for such varieties, pushing back against potentially discriminatory practices (in many cases, discriminatory through oversight more than malice). Our work calls for NLP researchers to consider both social and racial hierarchies sustained or intensified by current computational linguistic research. By shifting towards a human-in-the-loop paradigm to conduct deep multi-layered dialectal language analysis of AAE to counter-attack erasure and several forms of biases such as *selection bias*, *label bias*, *model over-amplification*, and *semantic bias* (see Shah et al. (2020) for definitions) in NLP.

We hope our dynamic approach can encourage practitioners, researchers and developers for AAE inclusive work, and that our contributions can pave the way for normalizing the use of a human-in-the-loop paradigm both to obtain new data and create NLP tools to better comprehend underrepresented dialect continua and English language varieties. In this way, NLP community can revolutionize the ways in which humans and technology cooperate by considering certain demographic attributes such as culture, background, race and gender when developing and deploying NLP models.

7 Limitations And Ethical Considerations

All authors must warrant that increased model performance for non-standard varieties such as underrepresented dialects, non-standard spellings or lexical items in NLP systems can potentially enable automated discrimination. In this work, we *solely* attempt to highlight the need for dialectal inclusivity for the development of impactful speech and language technologies in the future, and do not intend for increased feelings of marginalization of an already stigmatized community.

8 Acknowledgements

The authors would like to thank Shaylnn L.A. Crum-Dacon, Serena Lotreck, Brianna Brown and Kenia Segura Abá, Jyothi Kumar and Shin-Han Shiu for their support and the anonymous reviewers for their constructive comments.

¹³<https://oraal.uoregon.edu/coraal>

References

- Guy Bailey, John Baugh, Salikoko S. Mufwene, and John R. Rickford. 1998. *African-American English: Structure, History and Use (1st ed.)*. Routledge.
- John Baugh. 2008. Linguistic discrimination. In *1. Halbband*, pages 709–714. De Gruyter Mouton.
- Steven Bird. 2020. [Decolonising speech and language technology](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3504–3519, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Linda M. Bland-Stewart. 2005. Difference or deficit in speakers of african american english? <https://leader.pubs.asha.org/doi/10.1044/leader.FTR1.10062005.6>.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Su Lin Blodgett, Lisa Green, and Brendan O’Connor. 2016. [Demographic dialectal variation in social media: A case study of African-American English](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1119–1130, Austin, Texas. Association for Computational Linguistics.
- Su Lin Blodgett and Brendan O’Connor. 2017. [Racial disparity in natural language processing: A case study of social media african-american english](#). *CoRR*, abs/1707.00061.
- 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)*, pages 1415–1425, Melbourne, Australia. Association for Computational Linguistics.
- Herbert H. Clark and Michael F. Schober. 1992. Asking questions and influencing answers. In *Russell Sage Foundation*.
- Jamell Dacon and Haochen Liu. 2021. [Does gender matter in the news? detecting and examining gender bias in news articles](#). In *Companion Proceedings of the Web Conference 2021*, WWW ’21, page 385–392, New York, NY, USA. Association for Computing Machinery.
- Thomas Davidson and Debasmita Bhattacharya. 2020. [Examining racial bias in an online abuse corpus with structural topic modeling](#). *CoRR*, abs/2005.13041.
- Rachel Dorn. 2019. [Dialect-specific models for automatic speech recognition of African American Vernacular English](#). In *Proceedings of the Student Research Workshop Associated with RANLP 2019*, pages 16–20, Varna, Bulgaria. INCOMA Ltd.
- Yanai Elazar and Yoav Goldberg. 2018. [Adversarial removal of demographic attributes from text data](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 11–21, Brussels, Belgium. Association for Computational Linguistics.
- Anjalie Field, Su Lin Blodgett, Zeerak Waseem, and Yulia Tsvetkov. 2021. [A survey of race, racism, and anti-racism in NLP](#). *CoRR*, abs/2106.11410.
- Alex Graves and Jürgen Schmidhuber. 2005. [Framewise phoneme classification with bidirectional lstm and other neural network architectures](#). *Neural Networks*, 18(5):602–610. IJCNN 2005.
- Jonathon Green. 2014. *The vulgar tongue: Green’s history of slang*. Oxford University Press, New York, USA.
- Lisa J. Green. 2002. *African American English: A Linguistic Introduction*. Cambridge University Press.
- Sophie Groenwold, Lily Ou, Aesha Parekh, Samhita Honnavalli, Sharon Levy, Diba Mirza, and William Yang Wang. 2020. [Investigating African-American Vernacular English in transformer-based text generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5877–5883, Online. Association for Computational Linguistics.
- Matan Halevy, Camille Harris, Amy Bruckman, Diyi Yang, and Ayanna Howard. 2021. [Mitigating racial biases in toxic language detection with an equity-based ensemble framework](#). New York, NY, USA. Association for Computing Machinery.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Comput.*, 9(8):1735–1780.
- Taylor Jones. 2015. [Toward a description of african american vernacular english dialect regions using “black twitter”](#). *American Speech*, 90:403–440.
- Anna Jørgensen, Dirk Hovy, and Anders Søgaard. 2016. [Learning a POS tagger for AAVE-like language](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1115–1120, San Diego, California. Association for Computational Linguistics.
- Sharese King. 2020. [From african american vernacular english to african american language: Rethinking the study of race and language in african americans’ speech](#). *Annual Review of Linguistics*, 6(1):285–300.
- Allison Koenecke, Andrew Nam, Emily Lake, Joe Nudell, Minnie Quartey, Zion Mengesha, Connor Toups, John R. Rickford, Dan Jurafsky, and Sharad

- Goel. 2020. [Racial disparities in automated speech recognition](#). *Proceedings of the National Academy of Sciences*, 117(14):7684–7689.
- William Labov. 1975. [Ralph fasold, tense marking in black english: a linguistic and social analysis](#). Washington, d.c.: Center for applied linguistics, 1972. pp. 254. *Language in Society*, 4(2):222–227.
- John Lafferty, Andrew McCallum, and Fernando CN Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data.
- Wang Ling, Chris Dyer, Alan W Black, Isabel Trancoso, Ramón Fernández, Silvio Amir, Luís Marujo, and Tiago Luís. 2015. [Finding function in form: Compositional character models for open vocabulary word representation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1520–1530, Lisbon, Portugal. Association for Computational Linguistics.
- Haochen Liu, Jamell Dacon, Wenqi Fan, Hui Liu, Zitao Liu, and Jiliang Tang. 2020. [Does gender matter? towards fairness in dialogue systems](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4403–4416, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Edward Loper and Steven Bird. 2002. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL-02 Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics*, pages 63–70, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Marzieh Mozafari, Reza Farahbakhsh, and Noël Crespi. 2020. [Hate speech detection and racial bias mitigation in social media based on bert model](#). *PLOS ONE*, 15:1–26.
- Barbara Plank, Anders Søgaard, and Yoav Goldberg. 2016. [Multilingual part-of-speech tagging with bidirectional long short-term memory models and auxiliary loss](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 412–418, Berlin, Germany. Association for Computational Linguistics.
- 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*, pages 1668–1678, Florence, Italy. Association for Computational Linguistics.
- Deven Santosh Shah, H. Andrew Schwartz, and Dirk Hovy. 2020. [Predictive biases in natural language processing models: A conceptual framework and overview](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5248–5264, Online. Association for Computational Linguistics.
- Ian Stewart. 2014. [Now we stronger than ever: African-American English syntax in Twitter](#). In *Proceedings of the Student Research Workshop at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 31–37, Gothenburg, Sweden. Association for Computational Linguistics.
- Spencer S. Swinton. 1981. Predictive bias in graduate admissions tests. *ETS Research Report Series*, 1981.
- Rachael Tatman and Conner Kasten. 2017. [Effects of Talker Dialect, Gender & Race on Accuracy of Bing Speech and YouTube Automatic Captions](#). In *Proc. Interspeech 2017*, pages 934–938.
- Mengzhou Xia, Anjalie Field, and Yulia Tsvetkov. 2020. [Demoting racial bias in hate speech detection](#). In *Proceedings of the Eighth International Workshop on Natural Language Processing for Social Media*, pages 7–14, Online. Association for Computational Linguistics.
- Albert Xu, Eshaan Pathak, Eric Wallace, Suchin Gururangan, Maarten Sap, and Dan Klein. 2021. [Detoxifying language models risks marginalizing minority voices](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2390–2397, Online. Association for Computational Linguistics.
- Marcos Zampieri, Preslav Nakov, and Yves Scherrer. 2020. [Natural language processing for similar languages, varieties, and dialects: A survey](#). *Natural Language Engineering*, 26(6):595–612.
- Guanhua Zhang, Bing Bai, Junqi Zhang, Kun Bai, Conghui Zhu, and Tiejun Zhao. 2020. [Demographics should not be the reason of toxicity: Mitigating discrimination in text classifications with instance weighting](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4134–4145, Online. Association for Computational Linguistics.
- Xuhui Zhou, Maarten Sap, Swabha Swayamdipta, Yejin Choi, and Noah Smith. 2021. [Challenges in automated debiasing for toxic language detection](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3143–3155, Online. Association for Computational Linguistics.

A Dataset Details

Our collected dataset is demographically-aligned on AAE in correspondence on the dialectal tweet corpus by [Blodgett et al. \(2016\)](#). The Twitter-AAE corpus is publicly available and can be downloaded from link¹⁴. [Blodgett et al. \(2016\)](#) uses a

¹⁴<http://slanglab.cs.umass.edu/TwitterAAE/>

mixed-membership demographic language model which calculates demographic dialect proportions for a text accompanied by a race attribute—African America, Hispanic, Other, and White in that order. The race attribute is annotated by a jointly inferred probabilistic topic model based on the geolocation information of each user and tweet. Given that geolocation information (residence) is highly associated with the race of a user, the model can make accurate predictions. However, there are a low number of messages that possess a posterior probabilities of NaN as these are messages that have no in-vocabulary words under the model.

B Annotator Annotation Guidelines

You will be given demographically-aligned African American tweets, in which we refer to these tweets as sequences. As a dominant AAE speaker, who identifies as bi-dialectal, your task is to correctly identify the context of each word in a given sequence in hopes to address the issues of lexical, semantic and syntactic ambiguity.

1. Are you a dominant AAE speaker?
2. If you responded “yes” above, are you bi-dialectal?
3. If you responded “yes”, given a sequence, have you ever said, seen or used any of these words given the particular sequence?
4. Given a sequence, what are the SAE equivalents to the identified non-SAE terms?
5. For morphological and phonological (dialectal) purposes, are these particular words spelled how would you say or use them?
6. If you responded “no” above, can you provide a different spelling along with its SAE equivalent?

B.1 Annotation Protocol

1. What is the context of each word given the particular sequence?
2. Given NLTK’s Penn Treebank Tagset¹⁵, what is the most appropriate POS tag for each word in the given sequence?

¹⁵<https://www.guru99.com/pos-tagging-chunking-nltk.html>

B.2 Human evaluation of POS tags Protocol

1. Given the tagged sentence, are there any misclassified tags?
2. If you responded “yes” above, can you provide a different POS tag, and state why it is different?

C Variable Rules Examples

In this section we present a few examples of simple, deterministic phonological and morphological language features or *current* variable rules which highlight several regional varieties of AAE which typically attain misclassified POS tags. Please note that a more exhaustive list of these rules is still being constructed as this work is still ongoing. Below are a few variable cases (MAE → AAE), some of which may have been previously shown in Table 2:

1. Consonant (‘t’) deletion (Adverb case) : e.g. “*just*” → “*jus*”; “*must*” → “*mus*”
2. Contractive negative auxiliary verbs replacement: “*doesn’t*” → “*don’t*”
3. Contractive (‘re) loss: e.g. “*you’re*” → “*you*”; “*we’re*” → “*we*”
4. Copula deletion: Deletion of the verb “**be**” and its variants, namely “**is**” and “**are**” e.g. “*He is on his way*” → “*He on his way*”; “*You are right*” → “*You right*”
5. Homophonic word replacement (Pronoun case): e.g. “*you’re*” → “*your*”
6. Indefinite pronoun replacement: e.g. “*anyone*” → “*anybody*”;
7. Interdental fricative loss (Coordinating Conjunction case): e.g. “*this*” → “*dis*”; “*that*” → “*dat*”; “*the*” → “*da*”
8. Phrase reduction (present/ future tense) ⇒ word (Adverb case): e.g. “*what’s up*” → “*was-sup*”; “*fixing to*” → “*finna*”
9. Present tense possession replacement: e.g. “*John has two apples*” → “*John got two apples*”; “*The neighbors have a bigger pool*” → “*The neighbors got a bigger pool*”
10. Remote past “**been**” + completive (‘done’): “*I’ve already done that*” → “*I been done that*”

11. Remote past “*been*” + completive (‘*did*’):
“*She already did that*” → “*She been did that*”
12. Remote past “*been*” + Present tense possession replacement: “*I already have food*” → “*I been had food*”; “*You already have those shoes*” → “*You been got those shoes*”
13. Term-fragment deletion: e.g. “*brother*” → “*bro*”; “*sister*” → “*sis*”; “*your*” → “*ur*”; “*suppose*” → “*pose*”; “*more*” → “*mo*”
14. Term-fragment replacement: “*something*” → “*sumn*”; “*through*” → “*thru*”; “*for*” → “*fa*”; “*nothing*” → “*nun*”