

Racial Disparity in Natural Language Processing: A Case Study of Social Media African-American English

Su Lin Blodgett

University of Massachusetts Amherst
Amherst, MA
blodgett@cs.umass.edu

Brendan O'Connor

University of Massachusetts Amherst
Amherst, MA
brenocon@cs.umass.edu

ABSTRACT

We highlight an important frontier in algorithmic fairness: disparity in the quality of natural language processing algorithms when applied to language from authors of different social groups. For example, current systems sometimes analyze the language of females and minorities more poorly than they do of whites and males. We conduct an empirical analysis of racial disparity in language identification for tweets written in African-American English, and discuss implications of disparity in NLP.

1 INTRODUCTION: DISPARITY IN NLP

As machine learned algorithms govern more and more real-world outcomes, how to make them fair—and what that should mean—is of increasing concern. One strand of research, heavily represented at the FAT-ML series of workshops,¹ considers scenarios where a learning algorithm must make decisions about people, such as approving prospective applicants for employment, or deciding who should be the targets of police actions [5], and seeks to develop learners or algorithms whose decisions have only small differences in behavior between persons from different groups [4] or that satisfy other notions of fairness (e.g. [12, 13]).

Another recent strand of research has examined a complementary aspect of bias and fairness: *disparate accuracy in language analysis*. Linguistic production is a critically important form of human behavior, and a major class of artificial intelligence algorithms—natural language processing, or language technologies—may or may not fairly analyze language produced by different types of authors [7]. For example, Tatman [20] finds that YouTube autocaptioning has a higher word error rate for female speakers than for male speakers in videos. This has implications for downstream uses of language technology:

- Viewing: users who rely on autocaptioning have a harder time understanding what women are saying in videos, relative to what men are saying.
- Access: search systems are necessary for people to access information online, and for videos they may depend on indexing text recognized from the audio. Tatman’s results [20] imply that such a search system will fail to find information produced by female speakers more often than for male speakers.

This bias affects interests of the speakers—it is more difficult for their voices to be communicated to the world—as well as other users, who are deprived of information or opinions from females, or more generally, any social group whose language experiences lower accuracy of analysis by language technologies.

¹<http://www.fatml.org/>

Gender and dialect are well-known confounds in speech recognition, since they can implicate pitch, timbre, and the pronunciation of words (the phonetic level of language); domain adaptation is always a challenge and research continues on how to apply domain transfer to speech recognizers across dialects [15]. And more broadly, decades of research in the field of *sociolinguistics* has documented an extensive array of both social factors that affect how people produce language (e.g. community, geography, ethnicity), and how specifically language is affected (e.g. the lexicon, syntax, semantics). We might expect a minority teenager in school as well as a white middle-aged software engineer to both speak English, but they may exhibit variation in their pronunciation, word choice, slang, or even syntactic structures. Dialect communities often align with geographic and sociological factors, as language variation emerges within distinct social networks, or is affirmed as a marker of social identity.

Dialects pose a challenge to fairness in NLP, because they entail language variation that is *correlated* to social factors, and we believe there needs to be greater awareness of dialects among technologists using and building language technologies. In the rest of this paper, we focus on the dialect of African-American English as used on Twitter, which previous work [3, 9, 11] has established is very prevalent and sometimes quite different than mainstream American English. We analyze an African-American English Twitter corpus (from Blodgett et al. [3], described in §3), and analyze racial disparity in language identification, a crucial first step in any NLP application. Our previous work found that off-the-shelf tools display racial disparity—they tend to erroneously classify messages from African-Americans as non-English more often than those from whites. We extend this analysis from 200 to 20,000 tweets, finding that the disparity persists when controlling for message length (§4), and evaluate the racial disparity for several black-box commercial services. We conclude with a brief discussion (§5).

2 AFRICAN-AMERICAN ENGLISH AND SOCIAL MEDIA

We focus on language in social media, which is often informal and conversational. Social media NLP tools may be used for, say, sentiment analysis applications, which seek to measure opinions from online communities. But current NLP tools are typically trained on traditional written sources, which are quite different from social media language, and even more so from dialectal social media language. Not only does this imply social media NLP may be of lower accuracy, but since language can vary across social groups, any such measurements may be biased—incorrectly representing ideas and opinions from people who use non-standard language.

Specifically, we investigate dialectal language in publicly available Twitter data, focusing on African-American English (AAE), a dialect of American English spoken by millions of people across the United States [6, 14, 18]. AAE is a linguistic variety with defined syntactic-semantic, phonological, and lexical features, which have been the subject of a rich body of sociolinguistic literature. In addition to the linguistic characterization, reference to its speakers and their geographical location or speech communities is important, especially in light of the historical development of the dialect. Not all African-Americans speak AAE, and not all speakers of AAE are African-American; nevertheless, speakers of this variety have close ties with specific communities of African-Americans [6].

The phenomenon of “BlackTwitter” has been noted anecdotally; indeed, African-American and Hispanic minorities were markedly over-represented in the early years of the Twitter service (as well as younger people) relative to their representation in the American general population.² It is easy to find examples of non-Standard American English (SAE) language use, such as:

- (1) *he woke af smart af educated af daddy af coconut oil af GOALS AF & shares food af*
- (2) *Bored af den my phone finna die!!!*

The first example has low punctuation usage (there is an utterance boundary after every “af”), but more importantly, it displays a key syntactic feature of the AAE dialect, a null copula: “he woke” would be written, in Standard American English, as “he is woke” (meaning, politically aware). “af” is an online-specific term meaning “as f—.” The second example displays two more traditional AAE features: “den” is a spelling of “then” which follows a common phonological transform in AAE (initial “th” changing to a “d” sound: “dat,” “dis,” etc. are also common), and the word “finna” is an auxiliary verb, short for “fixing to,” which indicates an immediate future tense (“my phone is going to die very soon”); it is part of AAE’s rich verbal auxiliary system capable of encoding different temporal semantics than mainstream English [6].

3 DEMOGRAPHIC MIXED MEMBERSHIP MODEL FOR SOCIAL MEDIA

In order to test racial disparity in social media NLP, [3] collects a large-scale AAE corpus from Twitter, inferring soft demographic labels with a mixed-membership probabilistic model; we use this same corpus and method, briefly repeating the earlier description of the method. This approach to identifying AAE-like text makes use of the connection between speakers of AAE and African-American neighborhoods; we harvest a set of messages from Twitter, cross-referenced against U.S. Census demographics, and then analyze words against demographics with a mixed-membership probabilistic model. The data is a sample of millions of publicly posted geo-located Twitter messages (from the Decahose/Gardenhose stream [17]), most of which are sent on mobile phones, by authors in the U.S. in 2013.

For each message, we look up the U.S. Census blockgroup geographic area that the message was sent in, and use race and ethnicity information for each blockgroup from the Census’ 2013 American Community Survey, defining four covariates: percentages of the

²<http://www.pewinternet.org/fact-sheet/social-media/>

population that are non-Hispanic whites, non-Hispanic blacks, Hispanics (of any race), and (non-Hispanic) Asians. Finally, for each user u , we average the demographic values of all their messages in our dataset into a length-four vector $\pi_u^{(census)}$.

Given this set of messages and author-associated demographics, we infer statistical associations between language and demographics with a mixed membership probabilistic model. It directly associates each of the demographic variables with a topic; i.e. a unigram language model over the vocabulary. The model assumes an author’s mixture over the topics tends to be similar to their Census-associated demographic weights, and that every message has its own topic distribution. This allows for a single author to use different types of language in different messages, accommodating multidialectal authors. The message-level topic probabilities θ_m are drawn from an asymmetric Dirichlet centered on $\pi_u^{(census)}$, whose scalar concentration parameter α controls whether authors’ language is very similar to the demographic prior, or can have some deviation. A token t ’s latent topic z_t is drawn from θ_m , and the word itself is drawn from ϕ_{z_t} , the language model for the topic. Thus, the model learns demographically-aligned language models for each demographic category. Our previous work [3] verifies that its African-American language model learns linguistic attributes known in the sociolinguistics literature to be characteristic of AAE, in line with other work that has also verified the correspondence of geographical AA prevalence to AAE linguistic features on Twitter [10, 19].

This publicly available corpus contains 59.2 million tweets. We filter its messages to ones strongly associated with demographic groups; for example, for each message we infer the posterior proportion of its tokens that came from the African-American language model, which can be high either due to demographic prior, or from a message that uses many words exclusive to the AA language model (topic); these proportions are available in the released corpus. When we filter to messages with AA proportion greater than 0.8, this results in AAE-like text. We call these *AA-aligned* messages and we also select a set of white-aligned messages in the same way.³

4 BIAS IN NLP TOOLS

4.1 Language identification

Language identification, the task of classifying the major world language in which a message is written, is a crucial first step in almost any web or social media text processing pipeline. For example, in order to analyze the opinions of U.S. Twitter users, one might throw away all non-English messages before running an English sentiment analyzer. (Some of the coauthors of this paper have done this as a simple expedient step in the past.)

A variety of methods for language identification exist [8]; social media language identification is particularly challenging since messages are short and also use non-standard language [1]. In

³While Blodgett et al. verified that the AA-aligned tweets contain well-known features of AAE, we hesitate to call these “AAE” and “SAE” corpora, since technically speaking they are simply demographically correlated language models. The Census refers to the categories as “Black or African-American” and “White” (codes B03002E4 and B03002E3 in ACS 2013). And, while Hispanic- and Asian-associated language models of Blodgett et al.’s model are also of interest, we focus our analysis here on the African-American and White language models.

fact, a popular language identification system, *langid.py* [16], classifies both example messages in §2 as Danish with more than 99.9% confidence.

We take the perspective that since AAE is a dialect of American English, it ought to be classified as English for the task of major world language identification. We hypothesize that if a language identification tool is trained on standard English data, it may exhibit disparate performance on AA- versus white-aligned tweets. In particular, we wish to assess the *racial disparity accuracy difference*:

$$p(\text{correct} \mid \text{Wh}) - p(\text{correct} \mid \text{AA}) \quad (1)$$

From manual inspection of a sample of hundreds of messages, it appears that nearly all white-aligned and AA-aligned tweets are actually English, so accuracy is the same as proportion of English predictions by the classifier. A disparity of 0 indicates a language identifier that is fair across these classes. (An alternative measure is the ratio of accuracies, corresponding to Feldman et al.’s disparate impact measure [4].)

4.2 Experiments

We conduct an evaluation of four different off-the-shelf language identifiers, which are popular and straightforward for engineers to use when building applications:

- ***langid.py* (software)**: One of the most popular open source language identification tools, *langid.py* was originally trained on over 97 languages and evaluated on both traditional corpora and Twitter messages [16].
- **IBM Watson (API)**: The Watson Developer Cloud’s Language Translator service supports language identification of 62 languages.⁴
- **Microsoft Azure (API)**: Microsoft Azure’s Cognitive Services supports language identification of 120 languages.⁵
- **Twitter (metadata)**: The output of Twitter’s in-house identifier, whose predictions are included in a tweet’s metadata (from 2013, the time of data collection), which Twitter intends to “help developers more easily work with targeted subsets of Tweet collections.”⁶
- **Google (API, excluded)**: We attempted to test Google’s language detection service,⁷ but it returned a server error for every message we gave it to classify.

We queried the remote API systems in May 2017.

From manual inspection, we observed that longer tweets are significantly more likely to be correctly classified, which is a potential confound for a race disparity analysis, since the length distribution is different for each demographic group. To minimize this effect in our comparisons, we group messages into four bins (shown in Table 1) according to the number of words in the message. For each bin, we sampled 2,500 AA-aligned tweets and 2,500 white-aligned tweets, yielding a total of 20,000 messages across the two categories

⁴<https://www.ibm.com/watson/developercloud/doc/language-translator/index.html>

⁵<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/overview#language-detection>

⁶https://blog.twitter.com/developer/en_us/a/2013/introducing-new-metadata-for-tweets.html

⁷<https://cloud.google.com/translate/docs/detecting-language>

and four bins.⁸ We limited pre-processing of the messages to fixing of HTML escape characters and removal of URLs, keeping “noisy” features of social media text such as @-mentions, emojis, and hashtags. We then calculated, for each bin in each category, the number of messages predicted to be in English by each classifier. Accuracy results are shown in Table 1.⁹

As predicted, classifier accuracy does increase as message lengths increase; classifier accuracy is generally excellent for all messages containing at least 10 tokens. This result agrees with previous work finding short texts to be challenging to classify (e.g. [2]), since there are fewer features (e.g. character n-grams) to give evidence for the language used.¹⁰

However, the classifier results display a disparity in performance among messages of similar length; for all but one length bin under one classifier, accuracy on the white-aligned sample is higher than on the AA-aligned sample. The disparity in performance between AA- and white-aligned messages is greatest when messages are short; the gaps in performance for extremely short messages ranges across classifiers from 6.6% to 19.7%. This gap in performance is particularly critical as 41.7% of all AA-aligned messages in the corpus as a whole have 5 or fewer tokens.¹¹

5 DISCUSSION

Are these disparities substantively significant? It is easy to see how statistical bias could arise in downstream applications. For example, consider an analyst trying to look at major opinions about a product or political figure, with a sentiment analysis system that only gathers opinions from messages classified as English by Twitter. For messages length 5 or less, opinions from African-American speakers will be shown to be $1 - 54.0/73.7 = 27\%$ less frequent than they really are, relative to white opinions. Fortunately, the accuracy disparities are often only a few percentage points; nevertheless, it is important for practitioners to keep potential biases like these in mind.

One way forward to create less disparate NLP systems will be to use domain adaptation and other methods to extend algorithms to work on different distributions of data; for example, our demographic model’s predictions can be used to improve a language identifier, since the demographic language model’s posteriors accurately identify some cases of dialectal English [3]. In the context of speech recognition, Lehr et al. [15] pursue a joint modeling approach, learning pronunciation model parameters for AAE and SAE simultaneously.

One important issue may be the limitation of perspective of technologists versus users. In striking contrast to Twitter’s (historically) minority-heavy demographics, major U.S. tech companies are notorious for their low representation of African-Americans and Hispanics; for example, Facebook and Google report only 1%

⁸Due to a data processing error, there are 5 duplicates (19,995 unique tweets); we report on all 20,000 messages for simplicity.

⁹We have made the 20,000 messages publicly available at: <http://slanglab.cs.umass.edu/TwitterAAE/>

¹⁰A reviewer asked if length is used as a feature; we know that the open-source *langid.py* system does not (explicitly) use it.

¹¹For most (system,length) combinations, the accuracy difference is significant under a two-sided t-test ($p < .01$) except for two rows ($t \leq 5$, *langid.py*, $p = .03$) and ($10 < t \leq 15$, Twitter, $p = 0.5$). Accuracy rate standard errors range from 0.04% to 0.9% ($\approx \sqrt{acc(1 - acc)/2500}$).

		AA Acc.	WH Acc.	Diff.
<i>langid.py</i>	$t \leq 5$	68.0	70.8	2.8
	$5 < t \leq 10$	84.6	91.6	7.0
	$10 < t \leq 15$	93.0	98.0	5.0
	$t > 15$	96.2	99.8	3.6
IBM Watson	$t \leq 5$	62.8	77.9	15.1
	$5 < t \leq 10$	91.9	95.7	3.8
	$10 < t \leq 15$	96.4	99.0	2.6
	$t > 15$	98.0	99.6	1.6
Microsoft Azure	$t \leq 5$	87.6	94.2	6.6
	$5 < t \leq 10$	98.5	99.6	1.1
	$10 < t \leq 15$	99.6	99.9	0.3
	$t > 15$	99.5	99.9	0.4
Twitter	$t \leq 5$	54.0	73.7	19.7
	$5 < t \leq 10$	87.5	91.5	4.0
	$10 < t \leq 15$	95.7	96.0	0.3
	$t > 15$	98.5	95.1	-3.0

Table 1: Percent of the 2,500 tweets in each bin classified as English by each classifier; *Diff.* is the difference (disparity on an absolute scale) between the classifier accuracy on the AA-aligned and white-aligned samples. t is the message length for the bin.

of their tech employees are African-American,¹² as opposed to 13.3% in the overall U.S. population,¹³ and the population of computer science researchers in the U.S. has similarly low minority representation. It is of course one example of the ever-present challenge of software designers understanding how users use their software; in the context of language processing algorithms, such understanding must be grounded in an understanding of dialects and sociolinguistics.

REFERENCES

- [1] Timothy Baldwin, Paul Cook, Marco Lui, Andrew MacKinlay, and Li Wang. 2013. How Noisy Social Media Text, How Diffrent Social Media Sources?. In *International Joint Conference on Natural Language Processing*. 356–364.
- [2] Timothy Baldwin and Marco Lui. 2010. Language identification: The long and the short of the matter. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 229–237.
- [3] Su Lin Blodgett, Lisa Green, and Brendan O’Connor. 2016. Demographic Dialectal Variation in Social Media: A Case Study of African-American English. *Proceedings of EMNLP (2016)*.
- [4] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 259–268.
- [5] Sharad Goel, Maya Perelman, Ravi Shroff, and David Alan Sklansky. 2017. Combatting police discrimination in the age of big data. *New Criminal Law Review: In International and Interdisciplinary Journal* 20, 2 (2017), 181–232.
- [6] Lisa J. Green. 2002. *African American English: A Linguistic Introduction*. Cambridge University Press.
- [7] Dirk Hovy and L. Shannon Spruit. 2016. The Social Impact of Natural Language Processing. In *Proceedings of ACL*.
- [8] Baden Hughes, Timothy Baldwin, Steven Bird, Jeremy Nicholson, and Andrew MacKinlay. 2006. Reconsidering Language Identification for Written Language Resources. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC’06)*. European Language Resources Association (ELRA). <http://aclweb.org/anthology/L06-1274>
- [9] Taylor Jones. 2015. Toward a Description of African American Vernacular English Dialect Regions Using “Black Twitter”. *American Speech* 90, 4 (2015), 403–440.
- [10] Anna Jørgensen, Dirk Hovy, and Anders Søgaard. 2016. Learning a POS tagger for AAVE-like language. In *Proceedings of NAACL*. Association for Computational Linguistics.
- [11] Anna Katrine Jørgensen, Dirk Hovy, and Anders Søgaard. 2015. Challenges of studying and processing dialects in social media. In *Proceedings of the Workshop on Noisy User-generated Text*. 9–18.
- [12] Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2016. Rawlsian fairness for machine learning. *arXiv preprint arXiv:1610.09559* (2016).
- [13] Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. 2016. Fairness in Learning: Classic and Contextual Bandits. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 325–333. <http://papers.nips.cc/paper/6355-fairness-in-learning-classic-and-contextual-bandits.pdf>
- [14] William Labov. 1972. *Language in the inner city: Studies in the Black English vernacular*. Vol. 3. University of Pennsylvania Press.
- [15] Maider Lehr, Kyle Gorman, and Izhak Shafran. 2014. Discriminative pronunciation modeling for dialectal speech recognition. In *Proc. Interspeech*.
- [16] M. Lui and T. Baldwin. 2012. *langid.py: An Off-the-shelf Language Identification Tool*. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL 2012), Demo Session, Jeju, Republic of Korea*. <http://www.aclweb.org/anthology-new/P/P12/P12-3005.pdf>
- [17] Fred Morstatter, Jrgen Pfeffer, Huan Liu, and Kathleen Carley. 2013. Is the Sample Good Enough? Comparing Data from Twitter’s Streaming API with Twitter’s Firehose. In *International AAAI Conference on Weblogs and Social Media*. <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM13/paper/view/6071>
- [18] John Russell Rickford. 1999. *African American vernacular English: Features, evolution, educational implications*. Wiley-Blackwell.
- [19] 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*. Association for Computational Linguistics, Gothenburg, Sweden, 31–37. <http://www.aclweb.org/anthology/E14-3004>
- [20] 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*. Association for Computational Linguistics, Valencia, Spain, 53–59. <http://www.aclweb.org/anthology/W/W17/W17-1606>

¹²<https://newsroom.fb.com/news/2016/07/facebook-diversity-update-positive-hiring-trends-show-progress/>

¹³<https://www.google.com/diversity/>

¹³<https://www.census.gov/quickfacts/table/RHI225215/00>