

Demographic Dialectal Variation in Social Media: A Case Study of African-American English

Su Lin Blodgett[†] Lisa Green* Brendan O'Connor[†]

[†]College of Information and Computer Sciences *Department of Linguistics
University of Massachusetts Amherst

Abstract

Though dialectal language is increasingly abundant on social media, few resources exist for developing NLP tools to handle such language. We conduct a case study of dialectal language in online conversational text by investigating African-American English (AAE) on Twitter. We propose a distantly supervised model to identify AAE-like language from demographics associated with geo-located messages, and we verify that this language follows well-known AAE linguistic phenomena. In addition, we analyze the quality of existing language identification and dependency parsing tools on AAE-like text, demonstrating that they perform poorly on such text compared to text associated with white speakers. We also provide an ensemble classifier for language identification which eliminates this disparity and release a new corpus of tweets containing AAE-like language.

Data and software resources are available at:
<http://slanglab.cs.umass.edu/TwitterAAE>

1 Introduction

Owing to variation within a standard language, regional and social dialects exist within languages across the world. These varieties or dialects differ from the standard variety in syntax (sentence structure), phonology (sound structure), and the inventory of words and phrases (lexicon). 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.

As many of these dialects have traditionally existed primarily in oral contexts, they have historically been underrepresented in written sources. Consequently, NLP tools have been developed from text which aligns with mainstream languages. With the rise of social media, however, dialectal language is playing an increasingly prominent role in online conversational text, for which traditional NLP tools may be insufficient. This impacts many applications: for example, dialect speakers' opinions may be mischaracterized under social media sentiment analysis or omitted altogether (Hovy and Spruit, 2016). Since this data is now available, we seek to analyze current NLP challenges and extract dialectal language from online data.

Specifically, we investigate dialectal language in publicly available Twitter data, focusing on African-American English (AAE), a dialect of Standard American English (SAE) spoken by millions of people across the United States. 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 (Green, 2002). Due to its widespread use, established history in the sociolinguistic literature, and demographic associations, AAE provides an ideal starting point for the development of a statistical model that uncovers dialectal

language. In fact, its presence in social media is attracting increasing interest for natural language processing (Jørgensen et al., 2016) and sociolinguistic (Stewart, 2014; Eisenstein, 2015; Jones, 2015) research.¹ In this work we:

- Develop a method to identify *demographically-aligned* text and language from geo-located messages (§2), based on distant supervision of geographic census demographics through a statistical model that assumes a soft correlation between demographics and language.
- Validate our approach by verifying that text aligned with African-American demographics follows well-known phonological and syntactic properties of AAE, and document the previously unattested ways in which such text diverges from SAE (§3).
- Demonstrate *racial disparity* in the efficacy of NLP tools for language identification and dependency parsing—they perform poorly on this text, compared to text associated with white speakers (§4, §5).
- Improve language identification for U.S. online conversational text with a simple ensemble classifier using our demographically-based distant supervision method, aiming to eliminate racial disparity in accuracy rates (§4.2).
- Provide a corpus of 830,000 tweets aligned with African-American demographics.

2 Identifying AAE from Demographics

The presence of AAE in social media and the generation of resources of AAE-like text for NLP tasks has attracted recent interest in sociolinguistic and natural language processing research; Jones (2015) shows that nonstandard AAE orthography on Twitter aligns with historical patterns of African-American migration in the U.S., while Jørgensen et al. (2015) investigate to what extent it supports well-known sociolinguistics hypotheses about AAE.

¹Including a recent linguistics workshop: <http://linguistlaura.blogspot.co.uk/2016/06/using-twitter-for-linguistic-research.html>

Both, however, find AAE-like language on Twitter through keyword searches, which may not yield broad corpora reflective of general AAE use. More recently, Jørgensen et al. (2016) generated a large unlabeled corpus of text from hip-hop lyrics, subtitles from *The Wire* and *The Boondocks*, and tweets from a region of the southeast U.S. While this corpus does indeed capture a wide variety of language, we aim to discover AAE-like language by utilizing finer-grained, neighborhood-level demographics from across the country.

Our approach to identifying AAE-like text is to first harvest a set of messages from Twitter, cross-referenced against U.S. Census demographics (§2.1), then to analyze words against demographics with two alternative methods, a seedlist approach (§2.2) and a mixed-membership probabilistic model (§2.3).

2.1 Twitter and Census data

In order to create a corpus of demographically-associated dialectal language, we turn to Twitter, whose public messages contain large amounts of casual conversation and dialectal speech (Eisenstein, 2015). It is well-established that Twitter can be used to study both geographic dialectal varieties² and minority languages.³

Some methods exist to associate messages with authors' races; one possibility is to use birth record statistics to identify African-American-associated names, which has been used in (non-social media) social science studies (Sweeney, 2013; Bertrand and Mullainathan, 2003). However, metadata about authors is fairly limited on Twitter and most other social media services, and many supplied names are obviously not real.

Instead, we turn to geo-location and induce a distantly supervised mapping between authors and the demographics of the neighborhoods they live in (O'Connor et al., 2010; Eisenstein et al., 2011b; Stewart, 2014). We draw on a set of geo-located Twitter messages, most of which are sent on mobile phones, by authors in the U.S. in 2013. (These are selected from a general archive of the “Gardenhose/Decahose” sample stream of public Twitter messages.)

²For example, of American English (Huang et al., 2015; Doyle, 2014).

³For example, Lynn et al. (2015) develop POS corpora and taggers for Irish tweets; see also related work in §4.1.

ter messages (Morstatter et al., 2013)). Geo-located users are a particular sample of the userbase (Pavalanathan and Eisenstein, 2015), but we expect it is reasonable to compare users of different races within this group.

We look up the U.S. Census blockgroup geographic area that the message was sent in; blockgroups are one of the smallest geographic areas defined by the Census, typically containing a population of 600–3000 people. We use race and ethnicity information for each blockgroup from the Census’ 2013 American Community Survey, defining four covariates: percentages of the population that are non-Hispanic whites, non-Hispanic blacks, Hispanics (of any race), and Asian.⁴ Finally, for each user u , we average the demographic values of all their messages in our dataset into a length-four vector $\pi_u^{(census)}$. Under strong assumptions, this could be interpreted as the probability of which race the user is; we prefer to think of it as a rough proxy for likely demographics of the author and the neighborhood they live in.

Messages were filtered in order to focus on casual conversational text; we exclude tweets whose authors had 1000 or more followers, or that (a) contained 3 or more hashtags, (b) contained the strings “http”, “follow”, or “mention” (messages designed to generate followers), or (c) were retweeted (either containing the string “rt” or marked by Twitter’s metadata as re-tweeted).

Our initial Gardenhose/Decahose stream archive had 16 billion messages in 2013; 90 million were geo-located with coordinates that matched a U.S. Census blockgroup. 59.2 million tweets from 2.8 million users remained after pre-processing; each user is associated with a set of messages and averaged demographics $\pi_u^{(census)}$.

2.2 Direct Word-Demographic Analysis

Given a set of messages and demographics associated with their authors, a number of methods could be used to infer statistical associations between language and demographics.

Direct word-demographic analysis methods use the $\pi_u^{(census)}$ quantities to calculate statistics at the word level in a single pass. An intuitive approach is to calculate the *average demographics per word*.

⁴See appendix for additional details.

For a token in the corpus indexed by t (across the whole corpus), let $u(t)$ be the author of the message containing that token, and w_t be the word token. The average demographics of word type w is:⁵

$$\pi_w^{(softcount)} \equiv \frac{\sum_t 1\{w_t = w\} \pi_{u(t)}^{(census)}}{\sum_t 1\{w_t = w\}}$$

We find that terms with the highest $\pi_{w,AA}$ values (denoting high average African-American demographics of their authors’ locations) are very non-standard, while Stewart (2014) and Eisenstein (2013) find large $\pi_{w,AA}$ associated with certain AAE linguistic features.

One way to use the $\pi_{w,k}$ values to construct a corpus is through a seedlist approach. In early experiments, we constructed a corpus of 41,774 users (2.3 million messages) by first selecting the $n = 100$ highest- $\pi_{w,AA}$ terms occurring at least $m = 3000$ times across the data set, then collecting all tweets from frequent authors who have at least 10 tweets and frequently use these terms, defined as the case when at least $p = 20\%$ of their messages contain at least one of the seedlist terms. Unfortunately, the n, m, p thresholds are ad-hoc.


2.3 Mixed-Membership Demographic-Language Model

The direct word-demographics analysis gives useful validation that the demographic information may yield dialectal corpora, and the seedlist approach can assemble a set of users with heavy dialectal usage. However, the approach requires a number of ad-hoc thresholds, cannot capture authors who only occasionally use demographically-aligned language, and cannot differentiate language use at the message-level. To address these concerns, we develop a mixed-membership model for demographics and language use in social media.

The model directly associates each of the four demographic variables with a topic; i.e. a unigram language model over the vocabulary.⁶ The model assumes an author’s mixture over the topics tends to

⁵ $\pi_{w,k}$ has the flavor of “soft counts” in multinomial EM. By changing the denominator to $\sum_t \pi_{u(t)}^{(census)}$, it calculates a unigram language model that sums to one across the vocabulary. This hints at a more complete modeling approach (§2.3).

⁶To build the vocabulary, we select all words used by at least 20 different users, resulting in 191,873 unique words; other words are mapped to an out-of-vocabulary symbol.



$$\theta_m \sim \text{Dir}(\alpha\pi_u), \quad \phi \sim \text{Dir}(\beta/V)$$

$$z_t \sim \theta_m, \quad w_z \sim \phi_{z_t}$$

Figure 1: Mixed-membership model for users (u), messages (m) and tokens (t). Observed variables have a double lined border.

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 (Figure 1).

Thus the model learns *demographically-aligned* language models for each demographic category. The model is much more tightly constrained than a topic model—for example, if $\alpha \rightarrow \infty$, θ becomes fixed and the likelihood is concave as a function of ϕ —but it still has more joint learning than a direct calculation approach, since the inference of a messages’ topic memberships θ_m is affected not just by the Census priors, but also by the language used. A tweet written by an author in a highly AA neighborhood may be inferred to be non-AAE-aligned if it uses non-AAE-associated terms; as inference proceeds, this information is used to learn sharper language models.

We fit the model with collapsed Gibbs sampling (Griffiths and Steyvers, 2004) with repeated sample updates for each token t in the corpus,

$$p(z_t = k \mid w, z_{-t}) \propto \frac{N_{wk} + \beta/V}{N_k + \beta} \frac{N_{mk} + \alpha\pi_{uk}}{N_m + \alpha}$$

where N_{wk} is the number of tokens where word w occurs under topic $z = k$, N_{mk} is the number of tokens in the current message with topic k , etc.; all counts exclude the current t position. We observed

convergence of the log-likelihood within 100 to 200 iterations, and ran for 300 total.⁷ We average together count tables from the last 50 Gibbs samples for analysis of posterior topic memberships at the word, message, and user level; for example, the posterior probability a particular user u uses topic k , $P(z = k \mid u)$, can be calculated as the fraction of tokens with topic k within messages authored by u .

We considered α to be a fixed control parameter; setting it higher increases the correlations between $P(z = k \mid u)$ and $\pi_{u,k}^{(census)}$. We view the selection of α as an inherently difficult problem, since the correlation between race and AAE usage is already complicated and imperfect at the author-level, and census demographics allow only for rough associations. We set $\alpha = 10$ which yields posterior user-level correlations of $P(z = AA \mid u)$ against $\pi_{u,AA}$ to be approximately 0.8.

This model has broadly similar goals as non-latent, log-linear generative models of text that condition on document-level covariates (Monroe et al., 2008; Eisenstein et al., 2011a; Taddy, 2013). The formulation here has the advantage of fast inference with large vocabularies (since the partition function never has to be computed), and gives probabilistic admixture semantics at arbitrary levels of the data. This model is also related to topic models where the selection of θ conditions on covariates (Mimno and McCallum, 2008; Ramage et al., 2011; Roberts et al., 2013), though it is much simpler without full latent topic learning.

In early experiments, we used only two classes (AA and not AA), and found Spanish terms being included in the AA topic. Thus we turned to four race categories in order to better draw out non-AAE language. This removed Spanish terms from the AA topic; interestingly, they did not go to the Hispanic topic, but instead to Asian, along with other foreign languages. In fact, the correlation between users’ Census-derived proportions of Asian populations, versus this posterior topic’s proportions, is only 0.29, while the other three topics correlate to their respective Census priors in the range 0.83 to 0.87. This indicates the “Asian” topic actually functions as a background topic (at least in part). Better modeling of demographics and non-English

⁷Our straightforward single core implementation (in Julia) spends 80 seconds for each iteration over 586 million tokens.

language interactions is interesting potential future work.

By fitting the model to data, we can directly analyze unigram probabilities within the model parameters ϕ , but for other analyses, such as analyzing larger syntactic constructions and testing NLP tools, we require an explicit corpus of messages.

To generate a user-based AA-aligned corpus, we collected all tweets from users whose posterior probability of using AA-associated terms under the model was at least 80%, and generated a corresponding white-aligned corpus as well. In order to remove the effects of non-English languages, and given uncertainty about what the model learned in the Hispanic and Asian-aligned demographic topics, we focused only on AA- and white-aligned language by imposing the additional constraint that each user’s combined posterior proportion of Hispanic or Asian language was less than 5%. Our two resulting user corpora contain 830,000 and 7.3 million tweets, for which we are making their message IDs available for further research (in conformance with the Twitter API’s Terms of Service). In the rest of the work, we refer to these as the AA- and white-aligned corpora, respectively.

3 Linguistic Validation

Because validation by manual inspection of our AA-aligned text is impractical, we turn to the well-studied phonological and syntactic phenomena that traditionally distinguish AAE from SAE. We validate our model by reproducing these phenomena, and document a variety of other ways in which our AA-aligned text diverges from SAE.

3.1 Lexical-Level Variation

We begin by examining how much AA- and white-aligned lexical items diverge from a standard dictionary. We used SCOWL’s largest wordlist with level 1 variants as our dictionary, totaling 627,685 words.⁸

We calculated, for each word w in the model’s vocabulary, the ratio

$$r_k(w) = \frac{p(w|z = k)}{p(w|z \neq k)}$$

where the $p(\cdot|\cdot)$ probabilities are posterior inferences, derived from averaged Gibbs samples of the

⁸<http://wordlist.aspell.net/>

sufficient statistic count tables N_{wk} .

We selected heavily AA- and white-aligned words as those where $r_{AA}(w) \geq 2$ and $r_{white}(w) \geq 2$, respectively. We find that while 58.2% of heavily white-aligned words were not in our dictionary, fully 79.1% of heavily AA-aligned words were not. While a high number of out-of-dictionary lexical items is expected for Twitter data, this disparity suggests that the AA-aligned lexicon diverges from SAE more strongly than the white-aligned lexicon.

3.2 Internet-Specific Orthographic Variation

We performed an “open vocabulary” unigram analysis by ranking all words in the vocabulary by $r_{AA}(w)$ and browsed them and samples of their usage. Among the words with high r_{AA} , we observe a number of Internet-specific orthographic variations, which we separate into three types: abbreviations (e.g. *llh*, *kmsl*), shortenings (e.g. *dwn*, *dnt*), and spelling variations which do not correlate to the word’s pronunciation (e.g. *axx*, *bxtch*). These variations do not reflect features attested in the literature; rather, they appear to be purely orthographic variations highly specific to AAE-speaking communities online. They may highlight previously unknown linguistic phenomena; for example, we observe that *thoe* (SAE *though*) frequently appears in the role of a discourse marker instead of its standard SAE usage (e.g. *Girl Madison outfit THOE*). This new use of *though* as a discourse marker, which is difficult to observe using the SAE spelling amidst many instances of the SAE usage, is readily identifiable in examples containing the *thoe* variant. Thus, non-standard spellings provide valuable windows into a variety of linguistic phenomena.

In the next section, we turn to variations which do appear to arise from known phonological processes.

3.3 Phonological Variation

Many phonological features are closely associated with AAE (Green, 2002). While there is not a perfect correlation between orthographic variations and people’s pronunciations, Eisenstein (2013) shows that some genuine phonological phenomena, including a number of AAE features, are accurately reflected in orthographic variation on social media. We therefore validate our model by verifying that spellings reflecting known AAE phonological features align closely with the AA topic.

AAE	Ratio	SAE
sholl	1802.49	sure
iont	930.98	I don't
wea	870.45	where
talmbout	809.79	talking about
sumn	520.96	something

Table 1: Of 31 phonological variant words, top five by ratio $r_{AA}(w)$. SAE translations are shown for reference.

We selected 31 variants of SAE words from previous studies of AAE phonology on Twitter (Jørgensen et al., 2015; Jones, 2015). These variations display a range of attested AAE phonological features, such as derhotacization (e.g. *brotha*), deletion of initial *g* and *d* (e.g. *iont*), and realization of voiced *th* as *d* (e.g. *dey*) (Rickford, 1999).

Table 1 shows the top five of these words by their $r_{AA}(w)$ ratio. For 30 of the 31 words, $r \geq 1$, and for 13 words, $r \geq 100$, suggesting that our model strongly identifies words displaying AAE phonological features with the AA topic. The sole exception is the word *brotha*, which appears to have been adopted into general usage as its own lexical item.

3.4 Syntactic Variation

We further validate our model by verifying that it reproduces well-known AAE syntactic constructions, investigating three well-attested AAE aspectual or preverbal markers: habitual *be*, future *gone*, and completive *done* (Green, 2002). Table 2 shows examples of each construction.

To search for the constructions, we tagged the corpora using the ARK Twitter POS tagger (Gimpel et al., 2011; Owoputi et al., 2013),⁹ which Jørgensen et al. (2015) show has similar accuracy rates on both AAE and non-AAE tweets, unlike other POS taggers. We searched for each construction by searching for sequences of unigrams and POS tags characterizing the construction; e.g. for habitual *be* we searched for the sequences *O-be-V* and *O-b-V*. Non-standard spellings for the unigrams in the patterns were identified from the ranked analysis of §3.2.

We examined how a message’s likelihood of using each construction varies with the message’s posterior probability of AA. We split all messages into deciles based on the messages’ posterior probabili-

⁹Version 0.3.2: <http://www.cs.cmu.edu/~ark/TweetNLP/>

Construction	Example	Ratio
<i>O-be/b-V</i>	<i>I be tripping bruh</i>	11.94
<i>gone/gne/gon-V</i>	<i>Then she gon be single Af</i>	14.26
<i>done/dne-V</i>	<i>I done laughed so hard that I'm weak</i>	8.68

Table 2: AAE syntactic constructions and the ratios of their occurrences in the AA- vs. white-aligned corpora (§2.3).




Figure 2: Proportion of tweets containing AAE syntactic constructions by messages’ posterior probability of AA. On the x-axis, 0.1 refers to the decile [0, 0.1).

ity of AA. From each decile, we sampled 200,000 messages and calculated the proportion of messages containing the three syntactic constructions.

For all three constructions, we observed the clear pattern that as messages’ posterior probabilities of AA increase, so does their likelihood of containing the construction. Interestingly, for all three constructions, frequency of usage peaks at approximately the [0.7, 0.8) decile. One possible reason for the decline in higher deciles might be tendency of high-AA messages to be shorter; while the mean number of tokens per message across all deciles in our samples is 9.4, the means for the last two deciles are 8.6 and 7.1, respectively.

Given the important linguistic differences between our demographically-aligned subcorpora, we hypothesize that current NLP tools may behave differently. We investigate this hypothesis in §4 and §5.

4 Lang ID Tools on AAE

4.1 Evaluation of Existing Classifiers

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

	AAE	White-Aligned
<i>langid.py</i>	13.2%	7.6%
Twitter-1	8.4%	5.9%
Twitter-2	24.4%	17.6%

Table 3: Proportion of tweets in AA- and white-aligned corpora classified as non-English by different classifiers. (§4.1)

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.

Hughes et al. (2006) review language identification methods; social media language identification is challenging since messages are short, and also use non-standard and multiple (often related) languages (Baldwin et al., 2013). Researchers have sought to model code-switching in social media language (Rosner and Farrugia, 2007; Solorio and Liu, 2008; Maharjan et al., 2015; Zampieri et al., 2013; King and Abney, 2013), and recent workshops have focused on code-switching (Solorio et al., 2014) and general language identification (Zubiaga et al., 2014). For Arabic dialect classification, work has developed corpora in both traditional and Romanized script (Cotterell et al., 2014; Malmasi et al., 2015) and tools that use n-gram and morphological analysis to identify code-switching between dialects and with English (Elfardy et al., 2014).

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. Lui and Baldwin (2012) develop *langid.py*, one of the most popular open source language identification tools, training it on over 97 languages from texts including Wikipedia, and evaluating on both traditional corpora and Twitter messages. 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. Since language identifiers are typically based on character n-gram features, they may get confused by the types of lexical/orthographic divergences seen in §3. To evaluate this hypothesis, we compare the behavior of existing language identifiers on our subcorpora.

We test *langid.py* as well as the output of Twitter’s in-house identifier, whose predictions are included in a tweet’s metadata (from 2013, the time of data

collection); the latter may give a language code or a missing value (*unk* or an empty/null value). We record the proportion of non-English predictions by these systems; *Twitter-1* does not consider missing values to be a non-English prediction, and *Twitter-2* does.

We noticed emojis had seemingly unintended consequences on *langid.py*’s classifications, so removed all emojis by characters from the relevant Unicode ranges. We also removed @-mentions.

User-level analysis We begin by comparing the classifiers’ behavior on the AA- and white-aligned corpora. Of the AA-aligned tweets, 13.2% were classified by *langid.py* as non-English; in contrast, 7.6% of white-aligned tweets were classified as such. We observed similar disparities for *Twitter-1* and *Twitter-2*, illustrated in Table 3.

It turns out these “non-English” tweets are, for the most part, actually English. We sampled and annotated 50 tweets from the tweets classified as non-English by each run. Of these 300 tweets, only 3 could be unambiguously identified as written in a language other than English.

Message-level analysis We examine how a message’s likelihood of being classified as non-English varies with its posterior probability of AA. As in §3.4, we split all messages into deciles based on the messages’ posterior probability of AA, and predicted language identifications on 200,000 sampled messages from each decile.

For all three systems, the proportion of messages classified as non-English increases steadily as the messages’ posterior probabilities of AA increase. As before, we sampled and annotated from the tweets classified as non-English, sampling 50 tweets from each decile for each of the three systems. Of the 1500 sampled tweets, only 13 (~0.87%) could be unambiguously identified as being in a language other than English.

4.2 Adapting Language Identification for AAE

Natural language processing tools can be improved to better support dialects; for example, Jørgensen et al. (2016) use domain adaptation methods to improve POS tagging on AAE corpora. In this section, we contribute a fix to language identification to correctly identify AAE and other social media messages as English.




Figure 3: Proportion of tweets classified as non-English by messages’ posterior probability of AA. On the x-axis, 0.1 refers to the decile $[0, 0.1)$.

4.2.1 Ensemble Classifier

We observed that messages where our model infers a high probability of AAE, white-aligned, or “Hispanic”-aligned language almost always are written in English; therefore we construct a simple ensemble classifier by combining it with *langid.py*.

For a new message \vec{w} , we predict its demographic-language proportions $\hat{\theta}$ via posterior inference with our trained model, given a symmetric α prior over demographic-topic proportions (see appendix for details). The ensemble classifier, given a message, is as follows:

- Calculate *langid.py*’s prediction \hat{y} .
- If \hat{y} is English, accept it as English.
- If \hat{y} is non-English, and at least one of the message’s tokens are in demographic model’s vocabulary: Infer $\hat{\theta}$ and return English only if the combined AA, Hispanic, and white posterior probabilities are at least 0.9. Otherwise return the non-English \hat{y} decision.

Another way to view this method is that we are effectively training a system on an extended Twitter-specific English language corpus softly labeled by our system’s posterior inference; in this respect, it is related to efforts to collect new language-specific Twitter corpora (Bergsma et al., 2012) or minority language data from the web (Ghani et al., 2001).

4.2.2 Evaluation

Our analysis from §4.1 indicates that this method would correct erroneous false negatives for AAE

Message set	<i>langid.py</i>	Ensemble
High AA	80.1%	99.5%
High White	96.8%	99.9%
General	88.0%	93.4%

Table 4: Imputed recall of English messages in 2014 messages. For the *General* set these are an approximation; see text.

messages in the training set for the model. We further confirm this by testing the classifier on a sample of 2.2 million geolocated tweets sent in the U.S. in 2014, which are not in the training set.

In addition to performance on the entire sample, we examine our classifier’s performance on messages whose posterior probability of using AA- or white-associated terms was greater than 0.8 within the sample, which in this section we will call high AA and high white messages, respectively. Our classifier’s precision is high across the board, at 100% across manually annotated samples of 200 messages from each sample.¹⁰ Since we are concerned about the system’s overall recall, we impute recall (Table 4) by assuming that all high AA and high white messages are indeed English. Recall for *langid.py* alone is calculated by $\frac{n}{N}$, where n is the number of messages predicted to be English by *langid.py*, and N is the total number of messages in the set. (This is the complement of Table 3, except evaluated on the test set.) We estimate the ensemble’s recall as $\frac{n+m}{N}$, where $m = (n_{flip})P(\text{English} | \text{flip})$ is the expected number of correctly changed classifications (from non-English to English) by the ensemble and the second term is the precision (estimated as 1.0). We observe the baseline system has considerable difference in recall between the groups which is solved by the ensemble.

We also apply the same calculation to the general set of all 2.2 million messages; the baseline classifies 88% as English. This is a less accurate approximation of recall since we have observed a substantial presence of non-English messages. The ensemble classifies an additional 5.4% of the messages as English; since these are all (or nearly all) correct, this

¹⁰We annotated 600 messages as English, not English, or not applicable, from 200 sampled each from general, high AA, and high white messages. Ambiguous tweets which were too short (e.g. “Gm”) or contained only named entities (e.g. “Tennessee”) were excluded from the final calculations. The resulting samples have 197/197, 198/198, and 200/200 correct English classifications, respectively.

reflects at least a 5.4% gain to recall.

5 Dependency Parser Evaluation

Given the lexical and syntactic variation of AAE compared to SAE, we hypothesize that syntactic analysis tools also have differential accuracy. Jørgensen et al. (2015) demonstrate this for part-of-speech tagging, finding that SAE-trained taggers had disparate accuracy on AAE versus non-AAE tweets.

We assess a publicly available syntactic dependency parser on our AAE and white-aligned corpora. Syntactic parsing for tweets has received some research attention; Foster et al. (2011) create a corpus of constituent trees for English tweets, and Kong et al. (2014)’s *Tweeboparser* is trained on a Twitter corpus annotated with a customized unlabeled dependency formalism; since its data was uniformly sampled from tweets, we expect it may have low disparity between demographic groups.

We focus on widely used syntactic representations, testing the *SyntaxNet* neural network-based dependency parser (Andor et al., 2016),¹¹ which reports state-of-the-art results, including for web corpora. We evaluate it against a new manual annotation of 200 messages, 100 randomly sampled from each of the AA- and white-aligned corpora described in §2.3.

SyntaxNet outputs grammatical relations conforming to the Stanford Dependencies (SD) system (de Marneffe and Manning, 2008), which we used to annotate messages using *Brat*,¹² comparing to predicted parses for reference. Message order was randomized and demographic inferences were hidden from the annotator. To increase statistical power relative to annotation effort, we developed a partial annotation approach to only annotate edges for the root word of the first major sentence in a message. Generally, we found that that SD worked well as a descriptive formalism for tweets’ syntax; we describe handling of AAE and Internet-specific non-standard issues in the appendix. We evaluate labeled recall of the annotated edges for each message set:

Parser	AA	Wh.	Difference
SyntaxNet	64.0 (2.5)	80.4 (2.2)	16.3 (3.4)
CoreNLP	50.0 (2.7)	71.0 (2.5)	21.0 (3.7)

¹¹Using the publicly available *mcparsface* model: <https://github.com/tensorflow/models/tree/master/syntaxnet>

¹²<http://brat.nlplab.org/>

Bootstrapped standard errors (from 10,000 message resamplings) are in parentheses; differences are statistically significant ($p < 10^{-6}$ in both cases).

The white-aligned accuracy rate of 80.4% is broadly in line with previous work (compare to the parser’s unlabeled accuracy of 89% on English Web Treebank full annotations), but parse quality is much worse on AAE tweets at 64.0%. We test the Stanford CoreNLP neural network dependency parser (Chen and Manning, 2014) using the *english_SD* model that outputs this formalism;¹³ its disparity is worse. Soni et al. (2014) used a similar parser¹⁴ on Twitter text; our analysis suggests this approach may suffer from errors caused by the parser.

6 Discussion and Conclusion

We have presented a distantly supervised probabilistic model that employs demographic correlations of a dialect and its speaker communities to uncover dialectal language on Twitter. Our model can also close the gap between NLP tools’ performance on dialectal and standard text.

This represents a case study in dialect *identification*, *characterization*, and ultimately language technology *adaptation* for the dialect. In the case of AAE, dialect identification is greatly assisted since AAE speakers are strongly associated with a demographic group for which highly accurate governmental records (the U.S. Census) exist, which we leverage to help identify speaker communities. The notion of non-standard dialectal language implies that the dialect is underrepresented or underrecognized in some way, and thus should be inherently difficult to collect data on; and of course, *many* other language communities and groups are not necessarily officially recognized. An interesting direction for future research would be to combine distant supervision with unsupervised linguistic models to automatically uncover such underrecognized dialectal language.

Acknowledgments: We thank Jacob Eisenstein, Taylor Jones, Anna Jørgensen, Dirk Hovy, and the anonymous reviewers for discussion and feedback.

¹³*pos,depparse* options in version 2015-04-20, using tokenizations output by *SyntaxNet*.

¹⁴The older Stanford *englishPCFG* model with dependency transform (via pers. comm.).

References

- Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov, and Michael Collins. Globally normalized transition-based neural networks. *arXiv preprint arXiv:1603.06042*, 2016.
- Timothy Baldwin, Paul Cook, Marco Lui, Andrew MacKinlay, and Li Wang. How noisy social media text, how different social media sources? In *International Joint Conference on Natural Language Processing*, pages 356–364, 2013.
- Shane Bergsma, Paul McNamee, Mossaab Bagdouri, Clayton Fink, and Theresa Wilson. Language identification for creating language-specific Twitter collections. In *Proceedings of the Second Workshop on Language in Social Media*, pages 65–74. Association for Computational Linguistics, 2012.
- Marianne Bertrand and Sendhil Mullainathan. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. Technical report, National Bureau of Economic Research, 2003.
- Danqi Chen and Christopher Manning. A fast and accurate dependency parser using neural networks. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 740–750, Doha, Qatar, October 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/D14-1082>.
- Ryan Cotterell, Adithya Renduchintala, Naomi Saphra, and Chris Callison-Burch. An Algerian Arabic-French code-switched corpus. In *Workshop on Free/Open-Source Arabic Corpora and Corpora Processing Tools Workshop Programme*, page 34, 2014.
- M. C. de Marneffe and C. D. Manning. *Stanford typed dependencies manual*. Technical report, last revised April 2015 edition, 2008.
- Gabriel Doyle. Mapping dialectal variation by querying social media. In *Proceedings of EACL*, pages 98–106, 2014.
- Jacob Eisenstein. Phonological factors in social media writing. In *Proc. of the Workshop on Language Analysis in Social Media*, pages 11–19, 2013.
- Jacob Eisenstein. Identifying regional dialects in online social media. In C. Boberg, J. Nerbonne, and D. Watt, editors, *Handbook of Dialectology*. Wiley, 2015.
- Jacob Eisenstein, Amr Ahmed, and Eric P. Xing. Sparse additive generative models of text. In *Proceedings of ICML*, pages 1041–1048, 2011a.
- Jacob Eisenstein, Noah A. Smith, and Eric P. Xing. Discovering sociolinguistic associations with structured sparsity. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 1365–1374. Association for Computational Linguistics, 2011b.
- Heba Elfardy, Mohamed Al-Badrashiny, and Mona Diab. Aida: Identifying code switching in informal Arabic text. *Proceedings of EMNLP 2014*, page 94, 2014.
- Jennifer Foster, Ozlem Cetinoglu, Joachim Wagner, Joseph Le Roux, Stephen Hogan, Joakim Nivre, Deirdre Hogan, and Josef van Genabith. #hardtoparse: POS tagging and parsing the Twitterverse. In *Proc. of AAAI-11 Workshop on Analysing Microtext*, 2011.
- Rayid Ghani, Rosie Jones, and Dunja Mladenić. Mining the web to create minority language corpora. In *Proceedings of the Tenth International Conference on Information and Knowledge Management*, pages 279–286. ACM, 2001.
- Kevin Gimpel, Nathan Schneider, Brendan O’Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeff Flanigan, and Noah A. Smith. Part-of-speech tagging for Twitter: Annotation, features, and experiments. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 42–47. Association for Computational Linguistics, 2011.
- Lisa J. Green. *African American English: A Linguistic Introduction*. Cambridge University Press, 2002.
- T.L. Griffiths and M. Steyvers. Finding scientific topics. *Proceedings of the National Academy of*

- Sciences of the United States of America*, 101 (Suppl 1):5228, 2004.
- Dirk Hovy and L. Shannon Spruit. The social impact of natural language processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598. Association for Computational Linguistics, 2016. doi: 10.18653/v1/P16-2096. URL <http://aclweb.org/anthology/P16-2096>.
- Yuan Huang, Diansheng Guo, Alice Kasakoff, and Jack Grieve. Understanding US regional linguistic variation with Twitter data analysis. *Computers, Environment and Urban Systems*, 2015.
- Baden Hughes, Timothy Baldwin, Steven Bird, Jeremy Nicholson, and Andrew MacKinlay. 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), 2006. URL <http://aclweb.org/anthology/L06-1274>.
- Taylor Jones. Toward a description of African American Vernacular English dialect regions using “Black Twitter”. *American Speech*, 90(4): 403–440, 2015.
- Anna Jørgensen, Dirk Hovy, and Anders Søgaard. Learning a POS tagger for AAVE-like language. In *Proceedings of NAACL*. Association for Computational Linguistics, 2016.
- Anna Katrine Jørgensen, Dirk Hovy, and Anders Søgaard. Challenges of studying and processing dialects in social media. In *Proceedings of the Workshop on Noisy User-generated Text*, pages 9–18, 2015.
- Ben King and Steven P Abney. Labeling the languages of words in mixed-language documents using weakly supervised methods. In *Proceedings of HLT-NAACL*, pages 1110–1119, 2013.
- Lingpeng Kong, Nathan Schneider, Swabha Swayamdipta, Archna Bhatia, Chris Dyer, and Noah A. Smith. A dependency parser for tweets. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1001–1012, Doha, Qatar, October 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/D14-1108>.
- M. Lui and T. Baldwin. 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*, 2012. URL <http://www.aclweb.org/anthology-new/P/P12/P12-3005.pdf>.
- Teresa Lynn, Kevin Scannell, and Eimear Maguire. Minority language Twitter: Part-of-speech tagging and analysis of Irish tweets. *Proceedings of ACL-IJCNLP 2015*, page 1, 2015.
- Suraj Maharjan, Elizabeth Blair, Steven Bethard, and Tamar Solorio. Developing language-tagged corpora for code-switching tweets. In *The 9th Linguistic Annotation Workshop held in conjunction with NAACL 2015*, page 72, 2015.
- Shervin Malmasi, Eshrag Refaee, and Mark Dras. Arabic dialect identification using a parallel multidialectal corpus. In *International Conference of the Pacific Association for Computational Linguistics*, pages 35–53. Springer, 2015.
- David Mimno and Andrew McCallum. Topic models conditioned on arbitrary features with Dirichlet-Multinomial regression. In *Uncertainty in Artificial Intelligence*, pages 411–418, 2008.
- B. L. Monroe, M. P. Colaresi, and K. M. Quinn. Fightin’ Words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4):372, 2008.
- Fred Morstatter, Jrgen Pfeffer, Huan Liu, and Kathleen Carley. 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*, 2013. URL <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM13/paper/view/6071>.
- Brendan O’Connor, Jacob Eisenstein, Eric P. Xing, and Noah A. Smith. A mixture model of demographic lexical variation. In *NIPS Workshop on Machine Learning for Social Computing*, 2010.
- Olutobi Owoputi, Brendan O’Connor, Chris Dyer, Kevin Gimpel, Nathan Schneider, and Noah A. Smith. Improved part-of-speech tagging for on-

- line conversational text with word clusters. In *Proceedings of NAACL*, 2013.
- Umashanthi Pavalanathan and Jacob Eisenstein. Confounds and consequences in geotagged Twitter data. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*, Lisbon, September 2015. URL <http://www.aclweb.org/anthology/D/D15/D15-1256.pdf>.
- Daniel Ramage, Christopher D. Manning, and Susan Dumais. Partially labeled topic models for interpretable text mining. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 457–465, 2011.
- John Russell Rickford. *African American Vernacular English: Features, Evolution, Educational Implications*. Wiley-Blackwell, 1999.
- Margaret E Roberts, Brandon M Stewart, Dustin Tingley, and Edoardo M Airolidi. The structural topic model and applied social science. In *Advances in Neural Information Processing Systems Workshop on Topic Models: Computation, Application, and Evaluation*, 2013.
- Mike Rosner and Paulseph-John Farrugia. A tagging algorithm for mixed language identification in a noisy domain. In *Eighth Annual Conference of the International Speech Communication Association*, 2007.
- Thamar Solorio and Yang Liu. Learning to predict code-switching points. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 973–981. Association for Computational Linguistics, 2008.
- Thamar Solorio, Elizabeth Blair, Suraj Maharjan, Steven Bethard, Mona Diab, Mahmoud Ghoneim, Abdelati Hawwari, Fahad AlGhamdi, Julia Hirschberg, Alison Chang, and Pascale Fung. Overview for the first shared task on language identification in code-switched data. In *Proceedings of the First Workshop on Computational Approaches to Code Switching*, pages 62–72, Doha, Qatar, October 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/W14-3907>.
- Sandeep Soni, Tanushree Mitra, Eric Gilbert, and Jacob Eisenstein. Modeling factuality judgments in social media text. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 415–420, Baltimore, Maryland, June 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P14-2068>.
- Ian Stewart. Now we stronger than ever: African-American syntax in Twitter. *Proceedings of EACL*, page 31, 2014.
- Latanya Sweeney. Discrimination in online ad delivery. *ACM Queue*, 11(3):10, 2013.
- Matt Taddy. Multinomial inverse regression for text analysis. *Journal of the American Statistical Association*, 108(503):755–770, 2013.
- Marcos Zampieri, Binyam Gebrekidan Gebre, and Sascha Diwersy. N-gram language models and pos distribution for the identification of Spanish varieties. *Proceedings of TALN2013*, pages 580–587, 2013.
- Arkaitz Zubiaga, Inaki San Vicente, Pablo Gamallo, Jose Ramon Pichel, Inaki Algeria, Nora Aranberri, Aitzol Ezeiza, and Victor Fresno. Overview of TweetLID: Tweet language identification at SEPLN 2014. In *Proceedings of the Tweet Language Identification Workshop*, Girona, Spain, September 2014. Spanish Society for Natural Language Processing. URL <http://ceur-ws.org/Vol-1228/>.