Audience-Modulated Variation in Online Social Media

Umashanthi Pavalanathan and Jacob Eisenstein
School of Interactive Computing
Georgia Institute of Technology
{umashanthi+jacobe}@gatech.edu

Citation: American Speech (2015) 90(2): 187-213; doi:10.1215/00031283-3130324
http://americanspeech.dukejournals.org/content/90/2/187.full.pdf+html

Abstract

Stylistic variation in online social media writing is well attested: for example, geographical analysis of the social media service Twitter has replicated isoglosses for many known lexical variables from speech, while simultaneously revealing a wealth of new geographical lexical variables, including emoticons, phonetic spellings, and phrasal abbreviations. However, less is known about the *social* role of variation in online writing. This paper examines online writing variation in the context of audience design, focusing on affordances offered by Twitter that allow users to modulate a message's intended audience. We find that the frequency of non-standard lexical variables is inversely related to the size of the intended audience: as writers target smaller audiences, the frequency of lexical variables increases. In addition, these variables are more often used in messages that are addressed to individuals who are known to be geographically local. This phenomenon holds for geographically-differentiated lexical variables, but also for non-standard variables that are widely used throughout the United States. These findings suggest that users of social media are attuned to both the nature of their audience and the social meaning of lexical variation, and that they customize their self-presentation accordingly.

Introduction

Social media writing is often stylistically distinct from other written genres (Crystal 2006; Eisenstein 2013a), but it also displays an impressive internal stylistic diversity (Herring 2007; Androutsopoulos 2011). Many stylistic variables in social media have been shown to align with macro-level properties of the author, such as geographical location (Eisenstein et al. 2010), age (Schler et al. 2006), race (Eisenstein, Smith, and Xing 2011), and gender (Herring and Paolillo 2006). Linguistic differences are robust enough to support unnervingly accurate predictions of these characteristics based on writing style – with algorithmic predictions in some cases outperforming those of human judgments (Burger et al. 2011). This focus on prediction aligns with Silverstein's (2003) concept of first-order indexicality – the direct association of linguistic variables with macro-level social categories. The huge size of social media corpora makes it easy to identify hundreds of such variables through statistical analysis (e.g., Eisenstein, Smith, and Xing 2011).

But social media data has more to offer sociolinguistics than size alone: even though platforms such as Twitter are completely public, they capture language use in natural contexts with real social stakes. These platforms play host to a diverse array of interactional situations, from high school gossip to political debate, and from career networking to intense music fandom. As such, social media data offer new possibilities for understanding the social nature of language: not only who says what, but how stylistic variables are perceived by readers and writers, and how they are used to achieve communicative goals.

In this paper, we focus on the relevance of *audience* to sociolinguistic variation. A rich theoretical literature is already dedicated to this issue, including models of accommodation (Giles, Coupland, and Coupland 1991), audience design (Bell 1984), and stancetaking (Du Bois 2007). Empirical evidence for these models has typically focused on relatively small corpora of conversational speech, with a small number of hand-chosen variables. Indeed, the applicability of audience design and related models to a large-scale corpus of online written communication may appear doubtful – is audience a relevant and quantifiable concept in social media? In public "broadcast" media such as blogs, the properties of the audience seem difficult to identify. Conversely, in directed communication such as e-mails and SMS, the identity of the audience is clear, but acquisition of large amounts of data is impeded by obvious privacy considerations. However, ethnographic research suggests that users of Twitter have definite ideas about who their audience is, and that they customize their self-presentation accordingly (Marwick and boyd 2011). Furthermore, contemporary social media platforms such as Twitter and Facebook offer authors increasingly nuanced capabilities for manipulating the composition of their audience, enabling them to reach both within and beyond the social networks defined by explicitly-stated friendship ties (called "following" in Twitter; Kwak et al. 2010). We define these affordances in detail below.

This paper examines these notions of audience in the context of a novel dataset with thousands of writers and more than 200 lexical variables. The variables are organized into two sets: the first consists of terms that distinguish major American metropolitan areas from each other, and is obtained using an automatic technique based on regularized log-odds ratio. The second set of variables consists of the most frequently-used non-standard terms among Twitter users in the United States. In both cases, we find strong evidence of style-shifting according to audience size and proximity. When communication is intended for an individual recipient – particularly a recipient from the same geographical area as the author – both geographically-specific variables and medium-based variables are used at a significantly higher rate. Conversely, when communication is intended to reach a broad audience, outside the individual's social network, both types of variables are inhibited. These findings use a matched dataset design to control for the identify of the author, showing that individual authors are less likely to use non-standard and geographically-specific variables as the intended size of the audience grows. This provides evidence that individuals modulate their linguistic performance as they use social media affordances to control the intended audience of their messages. It also suggests that these non-standard variables – some of which appear to be endogenous to social media and recent in origin – are already viewed as socially marked, and are regulated accordingly.

Background

The main themes of this paper include language variation in social media and the creation of social meaning through style-shifting. We now review related work in these areas.

Language variation in social media

Social media comprises a range of Internet platforms, including collaborative writing projects such as Wikipedia, online communities such as Facebook and Myspace, forums such as Reddit and Stack Exchange, virtual game worlds, business and product reviews, and blog and microblogs (boyd and

Ellison 2007). These platforms offer many ways to interact with friends and strangers, but in the overwhelming majority of cases, interaction is conducted in written language. The relationship of social media writing to speech and to more traditional written media has been a topic of active interest, among both linguists (Walther and D'Addario 2001; Crystal 2006; Tagliamonte and Denis 2008; Dresner and Herring 2010) and computer scientists (Schler et al. 2006; Eisenstein 2013a), as well as in the popular press (Thurlow 2006). Computer scientists have been particularly interested in the relationship between language variation and macro-scale social variables. Formulated as a problem of predicting author attributes from text (Rao et al. 2010), they have demonstrated that it is possible to achieve high accuracy in predicting author age (Schler et al. 2006; Rosenthal and McKeown 2011; Nguyen et al. 2013), gender (Schler et al. 2006; Rao et al. 2010; Burger et al. 2011; Bamman, Eisenstein, and Schnoebelen 2014), race (Eisenstein, Smith, and Xing 2011), and geography (Cheng, Caverlee, and Lee 2010; Eisenstein et al. 2010). *Post hoc* analysis reveals that the most informative predictors include proper names (Wing and Baldridge 2011), spoken language dialect words (Doyle 2014), transcriptions of phonological variation (Eisenstein 2013b), as well as “netspeak” phenomena, such as emoticons and abbreviations (Eisenstein et al. 2012). A related line of research seeks linguistic predictors for social relationships such as power imbalances (Danescu-Niculescu-Mizil et al. 2012; Gilbert 2012; Prabhakaran, Rambow, and Diab 2012), and linguistic correlates for assessments of politeness (Danescu-Niculescu-Mizil et al. 2013). But while this work has identified and tested a wide range of sociolinguistic correlations, correlations alone cannot present the full picture of the social meaning of the associated variables (Eckert 1995; Johnstone and Kiesling 2008). In this paper, we use computational methods motivated by this prior work to identify geographically-oriented lexical variables, but we go further, showing how the distribution of these variables changes with the audience context.

Style-shifting and social meaning

The sociolinguistic interplay between macro-level speaker properties and interactional context has been a central concern of sociolinguistics since at least the 1960s, with Labov's 1966 study of the Lower East Side of New York City (Labov 2006) offering a classical example of *style-shifting*, in which the New York dialect is shown to be modulated by both socioeconomic status and the interactional situation. A key feature of this study is the interaction between fixed properties of the speaker (in this case, New York origin and socioeconomic status) and fluid properties of the interaction (in this case, attention to speech, which is experimentally modulated by the interviewer). By probing the linguistic consequences of this interaction between fixed and fluid social properties, we can arrive at an understanding of the social meaning of linguistic variation, which cannot be adduced from mere “first-order” associations between language and demographics (Eckert 1995; Silverstein 2003; Johnstone and Kiesling 2008) .

While Labov focused on attention to speech as an explanation for style-shifting, subsequent researchers have introduced other perspectives (Eckert and Rickford 2001). In communication accommodation theory, speakers adjust their linguistic style to converge or diverge, depending on their desire to reduce or increase social differences (Giles, Coupland, and Coupland 1991). The theory of audience design is related, but the target audience need not necessarily be the addressee (Bell 1984). This question of "which audience" is particularly salient in the context of publicly-readable social media: in principle any of the millions of users of Twitter could witness the conversation, though users may be aware of the specific preferences and interests of their own follower networks (Marwick and boyd 2011), and as we will see, can indirectly manipulate the composition of the audience through affordances offered by the platform.

Style-shifting in computer mediated communication

The literature on computer-mediated communication (CMC) has addressed the issue of style-shifting, although usually not in the large-scale empirical fashion considered in this paper. Androutsopoulos and Ziegler (2004) found evidence of regional language variation in the #mannheim Internet Relay Chat (IRC) channel, observing differences that match the North-South gradation of German dialects. They investigated style-shifting through manual analysis of a small number of example dialogues. More recently, Marwick and boyd (2011) conducted a series of interviews to understand how Twitter users navigate their "imagined audience", and the strategies they use to handle multiplicity of audience. They found that people consciously present themselves differently for different audiences on Twitter, by modulating their language, cultural references, and style. However, Bernstein et al. (2013) found that users of Facebook often have little idea of the size of the audience for their messages. Clearly, perceptions of audience vary across social media platforms, so the unique properties of each platform must be taken into account.

In a large scale study of Twitter, Danescu-Niculescu-Mizil, Gamon, and Dumais (2011) find evidence of linguistic accommodation in a large-scale corpus, but their linguistic variables mainly consist of lists of closed-class words such as articles and pronouns. They do not consider style-shifting in non-standard lexical varieties, as we do here. Paolillo (2001) examined linguistic variation on Internet Relay Chat (IRC) channel (#india) to test the hypothesis that standard linguistic variants tend to be associated with weak social ties, while vernacular variants are associated with strong network ties (Milroy and Milroy 1985). He used factor analysis to identify social characteristics of 94 members of the #india community, and then examined the use of five linguistic variables in conversations across members of different factor groups. The results were mixed, with some variables used more often among individuals with strong ties and other used more often among individuals with weak ties. This may be in part because the variables included phonetic spellings like *r* ("are") and *u* ("you"), which Paolillo notes were already widely accepted in computer-mediated communication in the 1990s. Our study focuses on linguistic variables that are not universally adopted, but rather, are strongly differentiated by geography. Our finding that non-standard variables are more likely to be used in messages that mention co-located individuals can be seen as consistent with Paolillo's original hypothesis, assuming ties to co-located individuals are usually stronger.

Paolillo (2001) also reports codeswitching between Hindi and English, and recent studies have replicated this finding in Facebook and Twitter. For example, Androutsopoulos (2014) investigate how a group of multilingual youth on Facebook strategically construct their audience using different language choices, finding that language choice is employed to maximize or partition the audience when starting new posts, or to align or disalign when responding to posts. Johnson (2013) studied a convenience sample of 25 Welsh/English bilingual users of Twitter, finding that they tend to use Welsh when writing to individuals who are also bilinguals, but write tweets in English when the message is not directed to any specific user. Nguyen et al. (2015) recently report similar results on a much larger scale, showing that Dutch Twitter users were more likely to employ the minority languages of Frisian and Limburgish in conversations with other users of these minority languages. Our study is related to work on code switching, but demonstrates audience-modulated variation in the use of English varieties.

The Social Environment of Twitter

Even though the number of potential readers for a public message in social media is nearly limitless, as Marwick and boyd (2011) suggest, social media users rely on a much more specific understanding of their audience as they make self-presentation decisions regarding language, cultural references and style. Our study focuses on Twitter, which provides a set of affordances that allow users to influence the likely composition of the audience for each message, shown in Figure 1. While any public message can in principle be read by anyone – we rely on this to build our dataset – these affordances can make it more or less likely that different types of people read any given message.

[Figure 1 about here]

- **Following**

  Twitter allows its users to "follow" other users: if User *A* follows User *B*, then most messages ("tweets") by User *B* will appear by default in the timeline shown on the main page of Twitter for User *A*. The exception is that messages specifically addressed to another User *C* only appear in the timeline of User *A* if she follows *both* Users *B* and *C*. Each user can view their follower list, and interviews conducted by Marwick and boyd (2011) indicate that users take their followers' expectations into account as they make decide how and what to write.

- **Hashtags**

  Prefixed by a "#" symbol with a keyword or phrase (e.g. #Ebola, #SierraLeone in Figure 1), a *hashtag* serves as a tagging convention on Twitter, enabling users to search outside their following network to find content of interest. Hashtags can be seen as creating a short-term virtual community of individuals who are interested on a particular topic or an event. Tweets with specific hashtags can be retrieved through Twitter Search, and people often search for hashtags to obtain tweets (either public or tweets from their followers) about a particular topic or an event (Chang 2010; Huang, Thornton, and Efthimiadis 2010). Messages that contain popular hashtags therefore have a higher likelihood of reaching an audience outside a user's follower set (Naaman, Becker, and Gravano 2011).

- **Usernames**

  Twitter messages can mention individual users by name, using the "@" symbol, e.g. @DeepakChopra in Figure 1; this convention dates back to Internet Relay Chat (IRC). In

> Twitter, when a username is the first token in a message, it can be viewed as a form of addressivity (Werry 1996). Honeycutt and Herring (2009) performed a quantitative study of the use of @-mentioning on Twitter and found that nearly 91% of the time it is used to direct a tweet to a specific addressee. The next most common use case was to refer another user, when the @-mention is in the body of the tweet, but not at the beginning. In Twitter, the use of this symbol gives the message special properties: by default, the individual who is @-mentioned receives a special notification, and if the message begins with an @-symbol, then it will not be visible by default to other individuals unless they follow both the sender and the recipient. For example, in Figure 1, the tweet from @Oprah to @DeepakChopra can be seen by users who follow both of them. Therefore, these messages are not completely private, and in the case of celebrities such as these individuals, it is probably best to view even these addressed messages as a public performance. But for the overwhelming majority of Twitter users, these messages will not be seen by anyone but the mentioned individual and perhaps a small number of mutual friends. The affordance can therefore be hypothesized to serve as narrowing the intended audience.

Following Kwak et al. (2010), we label messages that do not contain a hashtag or @-mention as *broadcasts*, since they are potentially viewed by all followers. Compared to broadcasts, messages that include hashtags are considered to target a larger audience, because broadcast messages are mostly viewed by only a user's followers while hashtag messages can reach a virtual community of users with common interests. The comparison between broadcasts, hashtags, and @-mentions allows us to quantify the effect of intended audience on linguistic style: we design predictors that capture each of these affordances, and test their effect on the use of two types of non-standard lexical items. In addition, within @-mentions, we differentiate pairs of users who are geolocated in the same metropolitan area.

Linguistic variables

Our analysis of audience-driven style-shifting focuses on non-standard lexical variables, which are easier to automatically quantify at scale than phonetic or morphosyntactic variables. We consider two sources of non-standard variables: lexical items that are strongly associated with specific geographical regions of the United States, and lexical items that are non-standard in other media, yet frequent in Twitter.

Geographical variables

Several papers have documented geographical lexical variation in Twitter (e.g., Eisenstein et al. 2010; see the related work section for more); as in this prior work, our approach is based on a large dataset of geotagged tweets, all from within the United States. We begin by geolocating each message's latitude and longitude coordinates to a Metropolitan Statistical Area (MSA), which is a geographical category defined by the U.S. Census Bureau as a high-density region organized around a single urban core. The largest MSA is New York-Newark-Jersey City (including parts of New York, New Jersey, and

Pennsylvania); the tenth largest is Boston-Cambridge-Newton (including parts of Massachusetts and New Hampshire).

Given a dataset of MSA-associated Twitter messages, the detection of geographical lexical variables can be seen as a problem of identifying words associated with document classes – an issue that has been examined in prior work (see Monroe, Colaresi, and Quinn 2008 for an overview). A straightforward approach would be to directly compare term frequencies, but this can be misleading. A difference-based comparison of term frequencies overstates the importance of common words, as a slight increase in the frequency of very common words such as *the* or *and* might be imperceptable (and could be due to exogenous factors such as the relative prevalence of English versus Spanish in a given area), yet it would dwarf the frequencies of more noticeable rare words, such as the classic sandwich variables 'hoagie' and 'grinder'. Conversely, a ratio-based comparison overstates the importance of rare words: in the limit, a word that appears just once in the entire dataset would have the maximum possible ratio of frequencies.

A solution proposed by Monroe, Colaresi, and Quinn (2008) is to compare the ratio of *regularized* log-odds, which has the effect of finding meaningful differences in the middle ground between very common and very rare words. We apply the non-parametric approach of Eisenstein, Ahmed, and Xing (2011), in which the probability of word w appearing in a document in group i is proportional to $\exp(m_w + \beta_{i,w})$, where the parameter m_w is shared across groups, and the parameter $\beta_{i,w}$ is shrunk towards zero through the application of a Laplace prior distribution, so that $\beta_{i,w} \sim \text{Laplace}(0,\lambda)$. This distribution is more sharply peaked at zero than the Gaussian distribution, and therefore aggressively shrinks the $\beta_{i,w}$ coefficients to zero, unless there is a strong difference in the word frequency across groups. By placing a non-parametric Jeffrey's prior on λ so that $p(\lambda) \propto 1/\lambda$, we can avoid the specification of any additional tuning parameters, while maintaining good generalization.

This statistical model is implemented in the SAGE software library,[i] which we use to obtain a list of the top 30 words (with the highest $\beta_{i,w}$ coefficients) for each of the ten most populous metropolitan statistical areas in the United States. We then manually removed proper names, such as *sixers* (a basketball team from Philadelphia) and *waffle* (the Waffle House franchise, popular in Atlanta), as well as non-English words and standard English words. All decisions were based on manual inspection of a randomly selected set of Tweets for each word; eliminated words are shown in the supplement. As shown in Table 1, our final list of linguistic variables contains 120 terms in total. Table 2 shows examples of tweets containing selected linguistic variables, and Figure 2 compares the national and local frequencies of each variable. A complete list is provided in the supplement.

[Table 1, Table 2, and Figure 2 around here]

'Tweetspeak' variables

In addition to the geographically-specific lexical variables, we create a list of variables which are broadly popular on Twitter, yet non-standard. This makes it possible to differentiate whether our results pertain specifically to geographical variation, or to non-standard language more generally. We begin by finding the 1000 most frequent terms in our sample of Twitter. We then automatically remove words that appear in a standard Unix dictionary, and manually remove terms that refer to entities, punctuation, special symbols, numbers, hashtags, and non-English words. Uncertain cases (e.g. 'y', which is a standard word in Spanish but a non-standard shortening of 'why' in English) were resolved by choosing the most frequent sense among twenty randomly-selected examples. Manually-removed cases are listed in the supplement. The final word list contains 94 non-standard terms, shown in Table 3.

[Table 3 around here]

Data

This study is performed on a dataset of social media text gathered from the public Gardenhose / Decahose version of the streaming API feed offered by the microblog site Twitter (Eisenstein et al. 2012). The dataset has been acquired by continuously pulling data from June 2009 to May 2012, and contains a total of 114 million geotagged messages from 2.77 million different user accounts; only messages geolocated within the United States were considered. Retweets – repetitions of previously posted messages – were eliminated using both Twitter metadata as well as the “RT” token (a common practice among Twitter users to indicate a retweet). Tweets containing URLs were eliminated so as to remove marketing-oriented messages, which are often automated. Accounts with more than 1000 followers or followees were removed for similar reasons. These filters are aggressive by design; they are intended to exploit the abundance of data in Twitter by focusing on a subset of messages that is highly likely to be originally written by the author. All messages were tokenized using the publicly-available Twokenize program[ii], and were downcased; no other textual preprocessing was performed.

Building a balanced corpus

To avoid potential confounds, we use resampling to balance the corpus, thus ensuring that each metropolitan area and each author contributes the same number of “positive” and “negative” messages, where positive messages contain a non-standard variable, and negative messages do not. This prevents larger geographical regions from overwhelming the analysis, and avoids more subtle author-level confounds. For example, suppose that some individuals habitually use Twitter in a more interactional manner (using more @-mentions), while at the same time choosing more non-standard linguistic variables – this behavior might intuitively be expected from young people. This would give the impression that @-mentions predict the use of non-standard lexical items, since both occur more frequently in the data from this subset of users; yet, for each individual user, the chance of using a non-standard variable might hypothetically be unaffected by whether the message contains an @-mention. By balancing the corpus so that each author contributes the same number of positive and negative messages, this potential confound is avoided.

To produce a balanced corpus, we begin by identifying the list of all tweets that contain each variable. From this list, we randomly sample a single message, and add it to our set of positive tweets. We then sample a message from the same author, requiring that it not contain any of the linguistic variables in our list. We repeat this process until there are a total of 1000 positive tweets for each variable. For eight variables, it was not possible to obtain 1000 examples, so we eliminated these from the analysis. The resulting dataset contains 224,000 tweets for the geographical variables, and 188,000 tweets for the Tweetspeak variables. Figure 1 in the supplement shows a histogram of the number of tweets per author in the dataset.

Geolocating message recipients

Each Twitter message contains GPS coordinates that we can geolocate to a Metropolitan Statistical Area (MSA). However, to determine whether the use of lexical variables is influenced by the location of the *target* of an @-message, we must also be able to geolocate usernames. Although many users report a geographical location in their Twitter profile, this may be inaccurate, or the place name may be written in a non-standard form (Hecht et al. 2011). Another possibility would be to link the username to a user ID, and then find the geographical coordinates of the corresponding tweets from that ID. While the username-ID mapping can be queried from Twitter, this is impractical for large-scale data, due to Twitter's policy of limiting the number of such queries per hour.

Instead, we linked usernames to the locations of the tweets that mention them (Jurgens 2013). If a username is mentioned by at least three different individuals within a given metropolitan area, and is never mentioned by anyone from outside that MSA, then we can guess with confidence that this MSA is the correct location for this username. Usernames which do not meet this criterion are treated as unknown. This threshold may be overly strict: a username must be popular enough to have three distinct mentioners, and can never be mentioned by anyone from outside the MSA. This criterion was chosen as a high-precision heuristic: while we undoubtedly fail to identify the location of some usernames, we can be fairly confident in the username-MSA mappings that are produced.

Analysis

Method

Our goal is to measure style-shifting by identifying audience-based factors that are statistically associated with the use of the linguistic variables we are interested in. We treat this as a binary prediction problem: in each tweet, was a linguistic variable used or not? This is a standard setting in variationist sociolinguistics, and we adopt the tool of logistic regression for this purpose (Tagliamonte 2006). Specifically, we treat the dependent variable as a draw from a Bernoulli distribution whose parameter arises from the logistic transformation of an inner product of predictors (features) and their associated weights. By analyzing the maximum-likelihood weights of relevant predictors, we can observe their impact on the use of non-standard linguistic variables in Twitter. We will specifically focus on predictors that characterize the intended audience of the message.

Predictors

Motivated by the different Twitter affordances that enable users to target messages to different audience, we design our predictors as follows:

- **Messages directed to wider audience:**

- **#-INIT:** *does the message begin with a hashtag?*
- **#-INTERNAL:** *does the message contain a hashtag, but not at the beginning?*

Hashtags (#) are used on Twitter as an index for messages that are about specific topics or events, and are therefore often targeted at virtual, ad hoc communities of users with similar interests. A positive weight for these predictors would suggest that non-standard linguistic variables are more likely to appear in messages that are intended for a broader audience than the follower network.

- **Messages directed to limited audience:**
  - **@-INIT:** *does the message begin with a username mention?*
  - **@-INTERNAL:** *does the message contain a username mention, but not at the beginning?*

Messages beginning with username mentions are often used as interactional replies, and are not visible by default to other followers of the author. Messages that mention a username internally are visible by default, but still create a special notification for the message recipient. Thus, a positive weight for these predictors would suggest that local linguistic variables are more likely to appear in messages that target limited audiences.

- **Messages directed to local audience**
  - **@-INIT-SAME-METRO:** *does the message begin with a username mention of an individual who is also located within the same metropolitan area?*
  - **@-INTERNAL-SAME-METRO:** *does the message mention a username of an individual who is also located within the same metropolitan area, but not at the beginning of the message?*

We specifically differentiate mentions of users who are in the same metropolitan area as the author. Such mentions could indicate particularly strong ties, or these messages could provide a context in which non-standard variables are used to claim covert local prestige (Trudgill 1972). Note that these predictors overlap with @-INIT and @-INTERNAL: messages tagged with @-INIT-SAME-METRO messages are a subset of those tagged with @-INIT. Hence, a positive weight for these predictors would suggest that they have additional power, indicating a special role for @-mentions between geographically co-located individuals.

In addition to these main experimental predictors, we employ a control for message length. Interactional messages tend to be shorter, and therefore afford fewer opportunities for the use of geographical linguistic variables. To control for this confound, we introduce message length predictors, with binary indicator predictors for the number of words in each message.

Models

These predictors are combined into two logistic regression models, summarized in Table 4. Both of these models test the frequency of non-standard variables against the baseline of broadcast messages, so that a positive coefficient for a given message type indicates greater tendency towards non-standard variables than in broadcast messages, and a negative coefficient indicates an inhibition of non-standard variables as compared with the same baseline.

[Table 4 around here]

Model-I

The first model tests the hypothesis that non-standard lexical variables are used more often in messages that target a limited audience. We therefore include the following predictors: @-INIT, @-INTERNAL, #-INIT, and #-INTERNAL. These predictors act as proxies for the size of the intended audience.

[Tables 5 and 6 around here]

Tables 5 and 6 show the results of logistic regression analysis, using Model-I predictors. The geographically-specific lexical variables are shown in Table 5, and the “Tweetspeak” variables are shown in Table 6. Results are broadly similar for both sets of variables. Both @-INIT and @-INTERNAL show strong positive association with the use of non-standard lexical variables, indicating that these variables are used more frequently in interactional contexts with more specifically-directed audiences. Conversely, the #-INIT and #-INTERNAL predictors have negative weights, suggesting that larger audiences inhibit the use of non-standard lexical variables. All of these predictor weights are statistically significant at $p < 0.01$.

Model-II

The second model tests the hypothesis that geographically-specific linguistic variables are used more frequently in messages targeted at individuals from the same geographical area. To capture this, we include the @-INIT-SAME-METRO and @-INTERNAL-SAME-METRO predictors, which are indicators for messages that reference a username that is geolocated to the same MSA as the message itself. These predictors overlap with @-INIT and @-INTERNAL, so if their coefficients are significantly greater than zero, this would indicate that non-standard lexical variables are even more likely to be used in messages targeting geographically-local users than in other interactional messages.

[Tables 7 and 8 around here]

Results are shown in Table 7 and 8, with geographically-specific lexical variables in Table 7, and the Tweetspeak variables are shown in Table 8. Results are again broadly similar for both sets of variables, with small but statistically significant positive coefficients for the new predictors. This indicates that non-standard lexical variables are especially likely to be used in messages that mention individuals from the same metropolitan area as the sender. While the coefficient for @-INTERNAL-SAME-METRO is slightly higher than the coefficient for @-INIT-SAME-METRO, the overlapping confidence intervals indicate that the difference is not statistically significant – although both are significantly greater than zero.

Discussion

Our results demonstrate a clear relationship between the use of audience-selection affordances in Twitter and the frequency of non-standard lexical variables. The direction of this relationship is consistent: as the audience becomes smaller and more local, non-standard variables are increasingly likely to be used; as the audience becomes larger and less directly connected to the author, the frequency of the variables decreases. We draw two main conclusions: first, that Twitter users are indeed aware of the audience-selecting role of affordances such as hashtags and username mentions, and second, that they are sensitive to the non-standardness of both the geographically-specific lexical variables as well as the Tweetspeak variables.

The sociolinguistics literature provides a range of theoretical accounts for style shifting, including situational variation (Finegan and Biber 1994), identity dimensions (Coupland 2001), and audience design (Bell 1984). We find that as the audience becomes larger, the frequency of non-standard variables decreases, and this aligns particularly well with the audience design framework of Bell, who notes how audience size affects the pressure to accommodate. In public speaking, the larger a speaker's audience, the greater the pressure to be understood and to win approval; Bell therefore posits a gradient from private to public situation of increasing influence by addressee on a speaker, with pressure to seek approval growing roughly with the size of the audience. In the case of Twitter, the hashtag feature can target of a larger audience, and we find the expected shift towards more standard language.

Coupland (2007) argues that audience design overemphasizes recipiency, while undervaluing individual identity construction; rather, stylistic variation should be seen as a dynamic presentation of speaker's identity (Coupland 2001). “Dialect style” features, which constitute regional differentiation, can be related in terms of stylistic strategies towards identity goals. This account would seem to cohere particularly well with our finding that the frequency of geographically-specific variables increases in messages addressed to individuals for whom the variable is local, so that these variables stake a claim towards local authenticity or identity for the author, in the context of a community in which these claims are particularly meaningful. However, this account does not explain the 'Tweetspeak' variables, which have no specific geographic orientation, yet are also used with higher frequency in messages addressed towards individuals from the same geographical region as the author. We must therefore seek other explanations. These variables can be seen as indexing other identities, such as youth culture or technical savvy (Androutsopoulos 2007). While such identities might be thought of as specifically non-local, it is nonetheless the case that strong social network ties – those ties that are most densely-embedded and most likely to be reciprocated – tend to be geographically local (McPherson, Smith-Lovin, and Cook 2001; McGee, Caverlee, and Cheng 2011). Thus, geographical co-location may be acting here as a proxy for tie strength, and it may be tie strength rather than geography that ultimately explains this generalized tendency towards non-standard language. As we have now reached the realm of speculation, we must conclude by noting that the rich social network information available in social media presents the opportunity for a new generation of research building on previous efforts to link tie strength to patterns of language variation and change (Milroy and Milroy 1985).

Finally, our analyses also have implications for the future of sociolinguistic diversity in social media.

Because social media services like Twitter permit rapid communication across geographical divides, it might be thought that they would contribute to leveling of dialect differences. However, our findings suggest the possibility of a countervailing force, as non-standard variables are most likely to be used in precisely those settings that are most similar to face-to-face communication: conversational messages aimed at narrow, local audiences. This tendency could help to explain why so much prior work has found that social media does not inhibit geographical language differentiation, but rather, seems to exacerbate it (Eisenstein et al. 2010; Wing and Baldridge 2011; Eisenstein 2015). A promising direction for future work would be to evaluate these network effects in a longitudinal context, which would connect this analysis with existing theories of how the social profile of linguistic variables change over time (Eckert 2000), and may also shed light on the implications of audience-modulated variation for the future prospects of diversity in online writing.

Tables

<table>
<tr><th>MSA</th><th>Linguistic Variables</th></tr>
<tr><td>New York</td><td>lml, deadass, od, odee, werd, cud, nuttin, nicee, sed, lata, buggin, wrd, noe, w|, layin, okk, lols, lolrt, crazyy, sour, wid</td></tr>
<tr><td>Los Angeles</td><td>fasho, ahah, cuh, koo, cuhz, fkn, ahahah, ;o</td></tr>
<tr><td>Chicago</td><td>mfs, goofy, nbs, lbvs, bogus, 2ma, lbs, mf, ikr, lmmfao, hoop, crackin</td></tr>
<tr><td>Dallas</td><td>ion, nun, oomf, tf, (;, finna, dang, fa, (:, <<, >>, <--, .!, trippin, y'all</td></tr>
<tr><td>Houston</td><td>mayne, fwm, jammin, shid, jamming, tripping, azz, bck, ma'am, bae, whoop, ole, sho, fck, lowkey, lawd, fa, trippin</td></tr>
<tr><td>Philadelphia</td><td>ard, jawn, cdfu, bul, wya, 1omf, jawns, ctfu, ctfuu, hbu, rd, foh, sike, hype, nut, bull</td></tr>
<tr><td>Washington DC</td><td>lt, lrt, llss, bait, fakin, stamp, ji, brova, siced, hu, wholetime, guh</td></tr>
<tr><td>Miami</td><td>bol, jit, bih, vibe</td></tr>
<tr><td>Atlanta</td><td>preciate, fye, frfr, slick, shid, fr, ain, ikr, followback, flex, gotcha</td></tr>
<tr><td>Boston</td><td>legit, deff, gunna</td></tr>
</table>

Table 1: *List of geographical linguistic variable from each ten most populous metropolitan statistical areas (MSAs) in the United States.*

(1) **lml** (love my life) thank u
(2) that was **od** (very) crazy
(3) dat (that) was **odee** (very) loud
(4) this line **deadass** (really) raps (wraps) around the whole store
(5) my sis **sed** (said) i wouldd lol
(6) o **werd** (really) ? lol i really like your style
(7) ima (I am a) la girl but these other ones be reaal **fkn** (fuking) ratchet man
(8) i'm so complicated smh (shake my head) **ion** (I don't) understand myself so yu (you) definitely can't understand me !!
(9) **finna** (fixing to) get ah (a) nap in and party lata (later)
(10) imm (I am) 2 hours late **fa** (for) skool (school).
(11) ohh **lawd** (Lord) i can't stand to be lied to !
(12) say wats (what's) wrong wit (with) **oomf** (one of my follower) she got life messed up .. #shitno
(13) lil (little) man **wya** (where you at) i'm **bouta** (about) come scoop u
(14) this **mf** (motherfucker) behind me is breathing so hard. i wanna **ctfu** (crack the fuck up) .
(15) im ready **fa** (for) college
(16) fa (for) **sho** (sure), hoe (derogatory term for woman, or promiscuous person of any gender)
(17) lol **jawns** (people) be lying
(18) bored as shyt (shit) ... **lls** (laughing like shit)
(19) yeah , yu (you) act like yu (you) don't **noe** (know) me !
(20) right lol . like **tf** (the fuck) is this bahaha
(21) **dang** (damn) didn't even mention ques .. hmm
(22) its copyrighted lol, **sike** (just kidding) u can use it
(23) #whatsthepoint of going to school tomorrow .? **lbvs** (laughing but very serious) .
(24) i **slick** (sort of, might) wanna move to memphis one day
(25) i'm **slick** (really) serious **af** (as fuck)
(26) these hoes (women) b (be) hatin hard **fr fr** (for real)
(27) oh **ard** (alright) just checkin lol
(28) i'm **gunna** (going to) punch the next stupid person that talks to me .
(29) i saw some stupid **shid** (shit) online today ..
(30) **preciate** (appreciate) that bro , the grey area coming very soon

Table 2: *Examples of some of the geographical linguistic variables in bold,with glosses in parentheses.*

lol	:)	dat	lmaoo	ohh	dm	luv
im	jus	dnt	lmfao	wats	<3	pics
lmao	ppl	idk	hmm	ahh	didnt	ii
ya	lil	aww	w/	nah	naw	comin
haha	wat	wtf	fuckin	umm	fam	tryin
nigga	yall	thats	talkin	outta	congrats	b4
da	yu	imma	tv	bday	yess	wassup
dont	omg	bro	:(	def	noo	;)
aint	goin	tryna	yay	lookin	oo	nothin
cuz	kno	hahaha	til	tha	watchin	abt
soo	ima	af	gon	shyt	alot	
mm	lmfaoo	txt	sayin	s/o	finna	
feelin	dats	gettin	dang	chillin	app	
smh	niggas	doin	pic	tht	bc	

Table 3: *List of Tweetspeak lexical variables*

Predictor	Model-I	Model-II	Description
Local audience features			
@-INIT-SAME-METRO		x	message begins with the username of an individual from the same MSA
@-INTERNAL-SAME-METRO		x	message contains a username of an individual from the same MSA, but not at the beginning
Limited audience features			
@-INIT	x	x	message begins with a username
@-INTERNAL	x	x	message contains a username, but not at the beginning
Wider audience features			
#-INIT	x	x	message begins with a hashtag
#-INTERNAL	x	x	message contains a hashtag, but not at the beginning
Controls			
NUM-WORDS-*i*	x	x	message contains exactly *i* tokens

Table 4: *Predictors used in each model*

Feature	Weight	Coefficient	95% C.I.	Empirical %	N
Limited audience					
@-INIT	0.5701	0.2821 ***	[0.264 , 0.300]	51.85	96,954
@-INTERNAL	0.5827	0.3340 ***	[0.299, 0.369]	56.41	15,494
Wider audience					
#-INIT	0.4004	-0.4037 ***	[-0.453, -0.355]	35.86	7,980
#-INTERNAL	0.4891	-0.0437 ***	[-0.076, -0.011]	50.40	16,937
Range	*18*				
Total				*50.00*	*224,000*

Table 5*: Results for Model-I predictors and geographical lexical variables. Statistical significance is indicated with asterisks,* *** *: p < 0.01,* ** *: p < 0.05,* * *: p > 0.05.* 'Weight' *is the logistic transformation of the logistic regression coefficient, yielding a value between 0 and 1, with 0.5 indicating indifference.* 'Empirical %' *indicates the percentage of messages with each predictor which include a non-standard variable, and* 'N' *is the total number of messages with each predictor.*

Feature	Weight	Coefficient	95% C.I.	Empirical %	N
Limited audience					
@-INIT	0.5997	0.4042 ***	[0.384, 0.425]	53.12	78,047
@-INTERNAL	0.5826	0.3333 ***	[0.294, 0.373]	54.12	12,076
Wider audience					
#-INIT	0.4079	-0.3728 ***	[-0.423,-0.323]	34.72	8,062
#-INTERNAL	0.4814	-0.0743 ***	[-0.108,-0.041]	49.10	16,472
Range	*19*				
Total				*50.00*	*188,000*

Table 6*: Results for Model-I predictors and Tweetspeak lexical variables.*

Feature	Weight	Coefficient	95% C.I.	Empirical %	N
Local audience					
@-INIT-SAME-METRO	0.5225	0.0900***	[0.055, 0.126]	53.23	14,976
@-INTERNAL-SAME-METRO	0.5272	0.1089**	[0.016, 0.202]	58.59	2,248
Limited audience					
@-INIT	0.5667	0.2685***	[0.249, 0.288]	51.85	96,954
@-INTERNAL	0.5789	0.3182***	[0.281, 0.355]	56.41	15,494
Wider audience					
#-INIT	0.4006	-0.4031***	[-0.452,-0.354]	35.86	7,980
#-INTERNAL	0.4894	-0.0424***	[-0.075,-0.010]	50.40	16,937
Range	*18*				
Total				*50.00*	*224,000*

Table 7: *Model-II predictors and geographical lexical variables.*

Feature	Weight	Coefficient	95% C.I.	Empirical %	N
Local audience					
@-INIT-SAME-METRO	0.5247	0.0990***	[0.050, 0.148]	53.64	7,349
@-INTERNAL-SAME-METRO	0.5523	0.2100***	[0.075, 0.345]	58.09	995
Limited audience					
@-INIT	0.5976	0.3954***	[0.375, 0.416]	53.12	78,047
@-INTERNAL	0.5783	0.3160***	[0.275, 0.357]	54.12	12,076
Wider audience					
#-INIT	0.4080	-0.3721***	[-0.422,-0.322]	34.72	8,062
#-INTERNAL	0.4816	-0.0735***	[-0.107,-0.040]	49.10	16,472
Range	*19*				
Total				*50.00*	*188,000*

Table 8: *Model-II predictors and Tweetspeak lexical variables.*

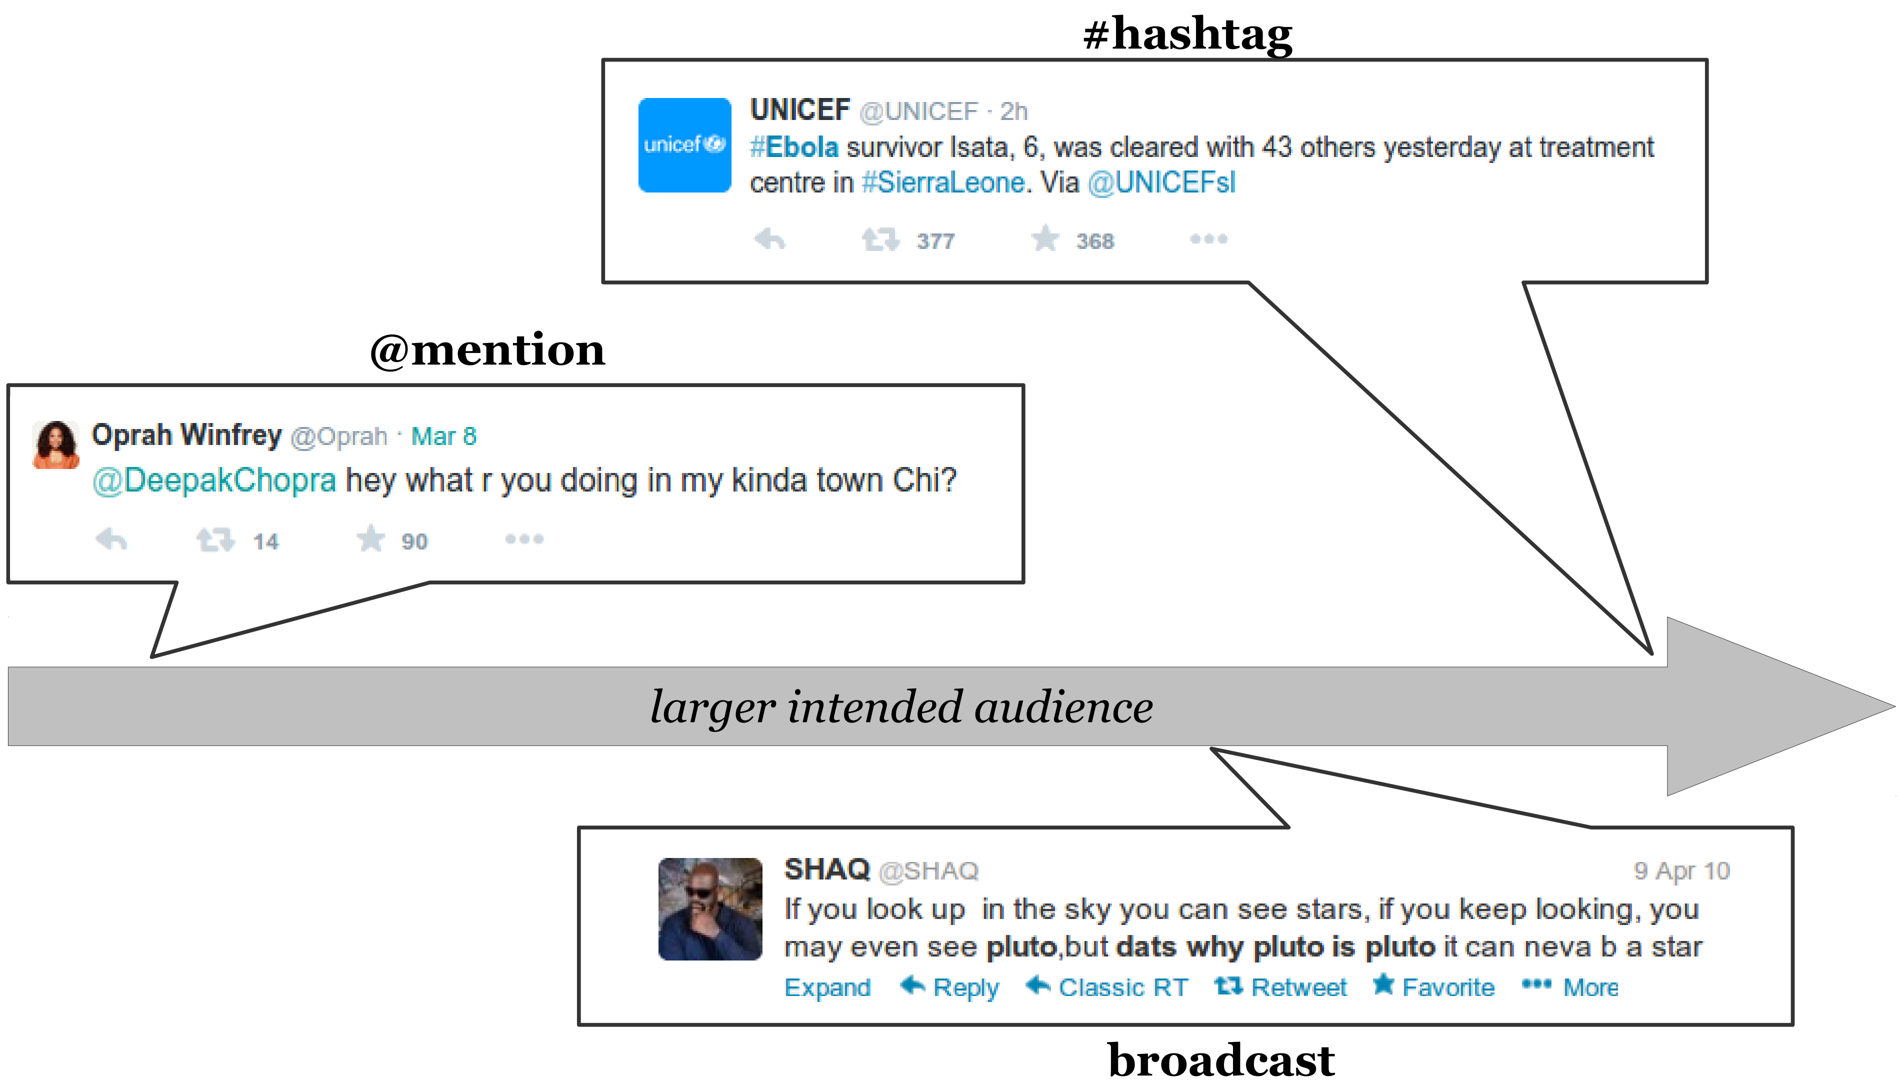
#hashtag
UNICEF @UNICEF · 2h
#Ebola survivor Isata, 6, was cleared with 43 others yesterday at treatment centre in #SierraLeone. Via @UNICEFsl
377
368
@mention
Oprah Winfrey @Oprah · Mar 8
@DeepakChopra hey what r you doing in my kinda town Chi?
14
90
larger intended audience
SHAQ @SHAQ
9 Apr 10
If you look up in the sky you can see stars, if you keep looking, you may even see pluto,but dats why pluto is pluto it can neva b a star
Expand
Reply
Classic RT
Retweet
Favorite
More
broadcast

Figure 1

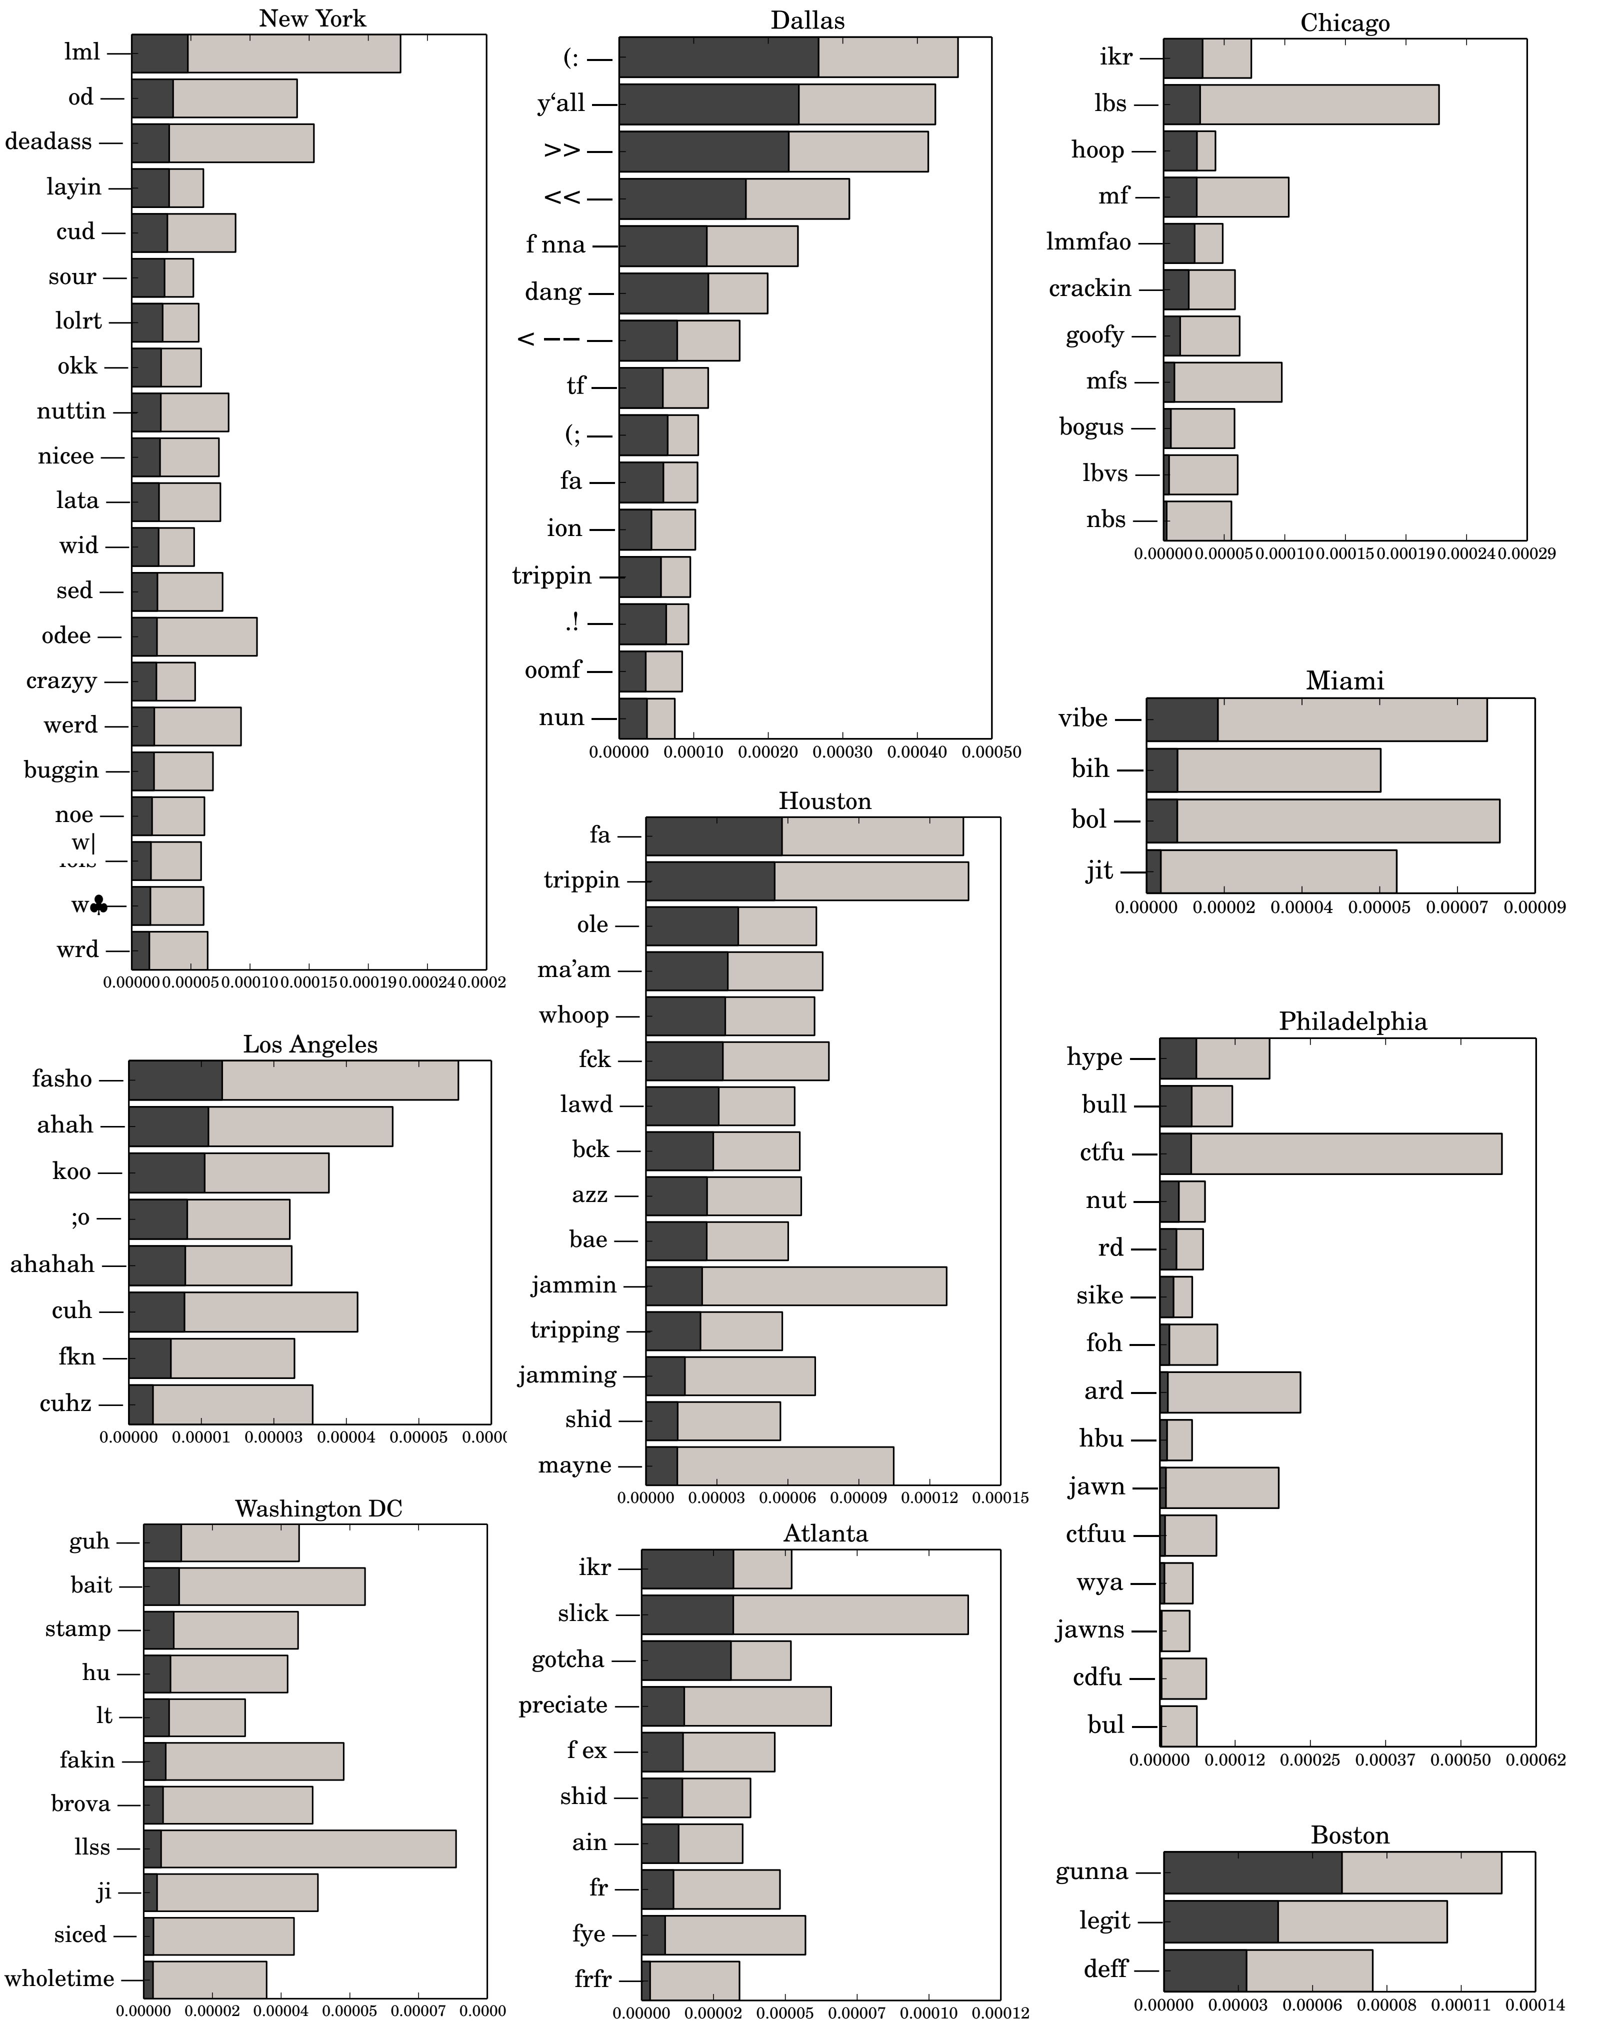
New York
lml
od
deadass
layin
cud
sour
lolrt
okk
nuttin
nicee
lata
wid
sed
odee
crazyy
werd
buggin
noe
wrd
Dallas
(:
y'all
f nna
dang
tf
(;
fa
ion
trippin
.!
oomf
nun
Chicago
ikr
lbs
hoop
mf
lmmfao
crackin
goofy
mfs
bogus
lbvs
nbs
Miami
vibe
bih
bol
jit
Los Angeles
fasho
ahah
koo
;o
ahahah
cuh
fkn
cuhz
Houston
fa
trippin
ole
ma'am
whoop
fck
lawd
bck
azz
bae
jammin
tripping
jamming
shid
mayne
Philadelphia
hype
bull
ctfu
nut
rd
sike
foh
ard
hbu
jawn
ctfuu
wya
jawns
cdfu
bul
Washington DC
guh
bait
stamp
hu
lt
fakin
brova
llss
ji
siced
wholetime
Atlanta
ikr
slick
gotcha
preciate
f ex
shid
ain
fr
fye
frfr
Boston
gunna
legit
deff

Figure 2

Figure captions

Figure 1: A*ffordances offered by Twitter for influencing the scope of the audience.*

Figure 2: *Frequencies of each linguistic variable from each metropolitan statistical area (MSA), with the local frequency in gray and the national frequency in black.*

i The specific software is available at https://github.com/jacobeisenstein/SAGE
ii The software is available at https://code.google.com/p/ark-tweet-nlp/

References

Androutsopoulos, Jannis. 2007. “Style online: Doing hip-hop on the German-speaking Web.” In *Style and social identities: alternative approaches to linguistic heterogeneity*, edited by Peter Auer, 279-317. Berlin, New York: de Gruyter.

———. 2011. “Language Change and Digital Media: A Review of Conceptions and Evidence.” In *Standard Languages and Language Standards in a Changing Europe*, edited by Kristiansen Tore and Nikolas Coupland, 145–61. Oslo: Novus.

———. 2014. “Languaging When Contexts Collapse: Audience Design in Social Networking.” *Discourse, Context & Media* 4: 62–73. Available at http://www.sciencedirect.com/science/article/pii/S2211695814000257.

Androutsopoulos, Jannis, and Evelyn Ziegler. 2004. “Exploring Language Variation on the Internet: Regional Speech in a Chat Community.” In *Language Variation in Europe: Papers from the Second International Conference on Language Variation in Europe 2,* edited by Britt-Louise Gunnarsson, 99–111. Uppsala: Uppsala University Press.

Bamman, David, Jacob Eisenstein, and Tyler Schnoebelen. 2014. “Gender Identity and Lexical Variation in Social Media.” *Journal of Sociolinguistics* 18 (02): 135–60. doi:10.1111/josl.12080.

Bell, Allan. 1984. “Language Style as Audience Design.” *Language in Society* 13 (2): 145–204. doi:10.1017/S004740450001037X.

Bernstein, Michael S, Eytan Bakshy, Moira Burke, and Brian Karrer. 2013. “Quantifying the Invisible Audience in Social Networks.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 27 April - 2 May 2013, Paris, France*, program chairs Stephen Brewster and Susanne Bødker, 21–30. New York: Association for Computing Machinery. doi:10.1145/2470654.2470658.

boyd, danah m., and Nicole B. Ellison. 2007. “Social Network Sites: Definition, History, and Scholarship.” *Journal of Computer-Mediated Communication* 13 (1): 210–30. doi:10.1111/j.1083-6101.2007.00393.x.

Burger, John D, John Henderson, George Kim, and Guido Zarrella. 2011. “Discriminating Gender on Twitter.” In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, 27–31 July 2011, Edinburgh, Scotland, United Kingdom*, program chairs Barzilay, Regina, and Mark Johnson, 1301–9. Pennsylvania: Association for Computational Linguistics.

Chang, Hsia-Ching. 2010. “A New Perspective on Twitter Hashtag Use: Diffusion of Innovation Theory.” *Proceedings of the American Society for Information Science and Technology* 47 (1): 1–4. New York: Wiley. doi:10.1002/meet.14504701295.

Cheng, Zhiyuan, James Caverlee, and Kyumin Lee. 2010. “You Are Where You Tweet: A Content-Based Approach to Geo-Locating Twitter Users.” In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management, 26–30 October 2010, Toronto, Canada*, program chairs Koudas, Nick, Gareth Jones, Xindong Wu, Kevyn Collins-Thompson, and Aijun An, 759–68. New York: Association for Computing Machinery. doi:10.1145/1871437.1871535.

Coupland, Nikolas. 2001. “Language, situation, and the relational self: Theorizing dialect-style in sociolinguistics.” In *Style and Sociolinguistic Variation*, edited by Eckert, Penelope, and John R Rickford, 185–210. Cambridge: Cambridge University Press.

———. 2007. *Style: Language Variation and Identity*. Cambridge: Cambridge University Press.

Crystal, David. 2006. *Language and the Internet*. Cambridge: Cambridge University Press.

Danescu-Niculescu-Mizil, Cristian, Michael Gamon, and Susan Dumais. 2011. "Mark My Words!: Linguistic Style Accommodation in Social Media." In *Proceedings of the 20th International Conference on World Wide Web, 28 March–1 April 2011, Hyderabad, India,* program chairs Elisa Bertino and Ravi Kumar, 745–54. New York: Association for Computing Machinery. doi:10.1145/1963405.1963509.

Danescu-Niculescu-Mizil, Cristian, Lillian Lee, Bo Pang, and Jon Kleinberg. 2012. "Echoes of Power: Language Effects and Power Differences in Social Interaction." In *Proceedings of the 21st International Conference on World Wide Web, 16–20 April 2012, Lyon, France,* program chairs Michael Rabinovich and Steffen Staab, 699–708. New York: Association for Computing Machinery. doi:10.1145/2187836.2187931.

Danescu-Niculescu-Mizil, Cristian, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. "A Computational Approach to Politeness with Application to Social Factors." In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, 4–9 August 2013, Sofia, Bulgaria*, program chairs Pascale Fung and Massimo Poesio, 250–59. Pennsylvania: Association for Computational Linguistics. http://aclweb.org/anthology/P13-1025.

Doyle, Gabriel. 2014. "A Computational Approach to Politeness with Application to Social Factors." In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, 26–30 April 2014, Gothenburg, Sweden*, program chairs Sharon Goldwater and Stefan Riezler, 98-106 . Pennsylvania: Association for Computational Linguistics. http://aclweb.org/anthology/E14-1011.

Dresner, Eli, and Susan C Herring. 2010. "Functions of the Nonverbal in CMC: Emoticons and Illocutionary Force." *Communication Theory* 20 (3): 249–68. New York: Wiley. doi:10.1111/j.1468-2885.2010.01362.x.

Du Bois, John W. 2007. "The Stance Triangle." In *Stancetaking in Discourse: Subjectivity, Evaluation, Interaction*, edited by Robert Englebretson, 139–82.

Eckert, Penelope. 1995. "(ay) Goes To the City. Exploring the Expressive Use of Variation." In *Amsterdam Studies in the Theory and History of Linguistic Science*, *Series 4*, edited by Guy, Gregory R., Crawford Feagin, Deborah Schiffrin and John Baugh, 47–68. Amsterdam/Pennsylvania: John Benjamins. doi:10.1075/cilt.127.06eck.

———. 2000. *Linguistic Variation as Social Practice: The Linguistic Construction of Social Meaning in Belten High*. Oxford: Blackwell.

Eckert, Penelope, and John R Rickford. 2001. *Style and Sociolinguistic Variation*. Cambridge: Cambridge University Press.

Eisenstein, Jacob. 2013a. "Phonological Factors in Social Media Writing." In *Proceedings of the Workshop on Language Analysis in Social Media,* edited by Cristian Danescu-Niculescu-Mizil, Atefeh Farzindar, Michael Gamon, Diana Inkpen, and Meena Nagarajan, *13 June 2013, Atlanta, Georgia, USA,* 11–19. Pennsylvania: Association for Computational Linguistics. http://aclweb.org/anthology/W13-1102.

———. 2013b. "What to Do about Bad Language on the Internet." In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 9–15 June 2013, Atlanta, Georgia, USA,* program chairs Hal Daume III and Katrin Kirchhoff, 359–69. Pennsylvania: Association for Computational Linguistics. http://aclweb.org/anthology/N13-1037.

———. 2015. "Written Dialect Variation in Online Social Media." In *The Handbook of*

Dialectology, edited by Boberg Charles, John Nerbonne, and Dominic Watt. New York: Wiley. (to appear).

Eisenstein, Jacob, Amr Ahmed, and Eric P Xing. 2011. "Sparse Additive Generative Models of Text." In *Proceedings of the 28th International Conference on Machine Learning (ICML-11), 28 June–2 July 2011, Bellevue, Washington*, program chairs Lise Getoor and Tobias Scheffer, 1041–48. Available at http://www.icml-2011.org/papers/534_icmlpaper.pdf.

Eisenstein, Jacob, Brendan O'Connor, Noah A Smith, and Eric P Xing. 2010. "A Latent Variable Model for Geographic Lexical Variation." In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, 9–11 October 2010, Massachusetts, USA*, program chairs Hang Li and Lluís Màrquez, 1277–87. Pennsylvania: Association for Computational Linguistics. http://dl.acm.org/citation.cfm?id=1870782.

———. 2012. "Mapping the Geographical Diffusion of New Words." *arXiv Preprint arXiv:1210.5268*. Available at http://arxiv.org/abs/1210.5268.

Eisenstein, Jacob, Noah A Smith, and Eric P Xing. 2011. "Discovering Sociolinguistic Associations with Structured Sparsity." In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1, 19–24 June 2011, Portland, Oregon, USA*, program chairs Yuji Matsumoto and Rada Mihalcea, 1365–74. Pennsylvania: Association for Computational Linguistics. http://dl.acm.org/citation.cfm?id=2002472.2002641.

Finegan, Edward, and Douglas Biber. 1994. "Register and social dialect variation: An integrated approach." In *Sociolinguistic Perspectives on Register*, edited by Douglas Biber and Edward Finegan, 315–347. Oxford: Oxford University Press.

Gilbert, Eric. 2012. "Phrases That Signal Workplace Hierarchy." In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work, 11–15 February 2012, Seattle, Washington, USA*, program chairs Jonathan Grudin, Gloria Mark, and John Riedl, 1037–46. New York: Association for Computing Machinery. doi:10.1145/2145204.2145359.

Giles, Howard, Justine Coupland, and Nikolas Coupland. 1991. *Contexts of Accommodation: Developments in Applied Sociolinguistics*. Cambridge: Cambridge University Press.

Hecht, Brent, Lichan Hong, Bongwon Suh, and Ed H Chi. 2011. "Tweets from Justin Bieber's Heart: The Dynamics of the Location Field in User Profiles." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 7–12 May 2011, Vancouver, British Columbia, Canada*, program chairs Geraldine Fitzpatrick, Carl Gutwin, Bo Begole, and Wendy A. Kellogg, 237–46. New York: Association for Computing Machinery. doi:10.1145/1978942.1978976.

Herring, Susan C. 2007. "A Faceted Classification Scheme for Computer-Mediated Discourse." *Language@ Internet* 4 (1): 1–37. http://www.languageatinternet.org/articles/2007/761 (accessed March 27, 2015).

Herring, Susan C, and John C Paolillo. 2006. "Gender and Genre Variation in Weblogs." *Journal of Sociolinguistics* 10 (4): 439–59. New York: Wiley. doi:10.1111/j.1467-9841.2006.00287.x.

Honey, Courtenay, and Susan C Herring. 2009. "Beyond Microblogging: Conversation and Collaboration via Twitter." In *System Sciences, 2009. HICSS'09. 42nd Hawaii International Conference on System Sciences, 5–8 January 2009, Big Island, Hawaii*, program chair Ralph H. Sprague Jr., 1–10. New Jersey: Institute of Electrical and Electronics Engineers. doi:10.1109/HICSS.2009.89.

Huang, Jeff, Katherine M Thornton, and Efthimis N Efthimiadis. 2010. "Conversational Tagging in Twitter." In *Proceedings of the 21st ACM Conference on Hypertext and Hypermedia, 13–16 June 2010, Toronto, Canada*, program chair Elaine Toms, 173–78. New York: Association for

Computing Machinery. http://dl.acm.org/citation.cfm?id=1810647.

Johnson, Ian. 2013. “Audience Design and Communication Accommodation Theory: Use of Twitter by Welsh-English Biliterates.” In *Social Media and Minority Languages: Convergence and the Creative Industries*, edited by Elin Haf Gruffydd Jones and Enrique Uribe-Jongbloed, 99–118. Bristol: Multilingual Matters.

Johnstone, Barbara, and Scott F Kiesling. 2008. “Indexicality and Experience: Exploring the Meanings Of/aw/-Monophthongization in Pittsburgh1.” *Journal of Sociolinguistics* 12 (1): 5–33. New York: Wiley. doi:10.1111/j.1467-9841.2008.00351.x.

Jurgens, David. 2013. “That’s What Friends Are For: Inferring Location in Online Social Media Platforms Based on Social Relationships.” In P*roceedings of the AAAI International Conference on Weblogs and Social Media, 8–10 July 2013, Boston, USA*, program chairs Nicole B. Ellison, Bernie Hogan, Paul Resnick, and Ian Soboroff, 273-82 . California: Association for the Advancement of Artificial Intelligence. Available at http://www.aaai.org/ocs/index.php/ICWSM/ICWSM13/paper/view/6067.

Kwak, Haewoon, Changhyun Lee, Hosung Park, and Sue Moon. 2010. “What Is Twitter, a Social Network or a News Media?” In *Proceedings of the 19th International Conference on World Wide Web, 26–30 April 2010, Raleigh, North Carolina, USA*, program chair Juliana Freire and Soumen Chakrabarti, 591–600. New York: Association for Computing Machinery. http://dl.acm.org/citation.cfm?id=1772751.

Labov, William. 1972. “The isolation of contextual styles.” *Sociolinguistic Patterns* : 70–100. Philadelphia: University of Pennsylvania Press.

———. 2006. *The social stratification of English in New York city*. Cambridge: Cambridge University Press.

Marwick, Alice E., and danah boyd. 2011. “I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience.” *New Media & Society* 13 (1): 114–33. doi:10.1177/1461444810365313.

McGee, Jeffrey, James A. Caverlee, and Zhiyuan Cheng. 2011. "A geographic study of tie strength in social media." *In Proceedings of the 20th ACM international conference on Information and knowledge management, 24–28 October 2011, Glasgow, Scotland, United Kingdom,* edited by Berendt, Bettina,vArjen de Vries, and Wenfei Fan, 2333–2336. New York: Association for Computing Machinery. http://dl.acm.org/citation.cfm?id=2063959.

McPherson, Miller, Lynn Smith-Lovin, and James M Cook. 2001. “Birds of a Feather: Homophily in Social Networks.” *Annual Review of Sociology* 27: 415–44. Palo Alto, Calif.:Annual Reviews. http://www.jstor.org/stable/2678628.

Milroy, James, and Lesley Milroy. 1985. “Linguistic change, social network and speaker innovation.” *Journal of Linguistics* 21 (02): 339–84. doi:10.1017/S0022226700010306.

Monroe, Burt L, Michael P Colaresi, and Kevin M Quinn. 2008. “Fightin’words: Lexical Feature Selection and Evaluation for Identifying the Content of Political Conflict.” *Political Analysis* 16 (4): 372–403. doi:10.1093/pan/mpn018.

Naaman, Mor, Hila Becker, and Luis Gravano. 2011. “Hip and Trendy: Characterizing Emerging Trends on Twitter.” *Journal of the American Society for Information Science and Technology* 62 (5): 902–18. doi: 10.1002/asi.21489.

Nguyen, Dong, Rilana Gravel, Dolf Trieschnigg, and Theo Meder. 2013. “‘How Old Do You Think I Am?’ A Study of Language and Age in Twitter.” In P*roceedings of the AAAI International*

Conference on Weblogs and Social Media, 8–10 July 2013, Boston, USA, program chairs Nicole B. Ellison, Bernie Hogan, Paul Resnick, and Ian Soboroff, 439–448. California: Association for the Advancement of Artificial Intelligence. Available at http://www.aaai.org/ocs/index.php/ICWSM/ICWSM13/paper/view/5984/6383.

Nguyen, Dong, Dolf Trieschnigg, and Leonie Cornips. 2015. "Audience and the Use of Minority Languages on Twitter." In *Proceedings of the AAAI International Conference on Web and Social Media, 26–29 May 2015, Boston, USA,* program chairs Meeyoung Cha, Cecilia Mascolo, and Christian Sandvig , California: Association for the Advancement of Artificial Intelligence. Available at http://www.dongnguyen.nl/publications/nguyen-icwsm2015.pdf. (accessed March 27, 2015).

Paolillo, John C. 2001. "Language Variation on Internet Relay Chat: A Social Network Approach." *Journal of Sociolinguistics* 5 (2): 180–213. New York: Wiley. doi:10.1111/1467-9481.00147.

Prabhakaran, Vinodkumar, Owen Rambow, and Mona Diab. 2012. "Predicting Overt Display of Power in Written Dialogs." In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, *3–8 June 2012, Montreal, Canada,* program chairs Eric Fosler-Lussier, Ellen Riloff, and Srinivas Bangalore, 518–22. Pennsylvania: Association for Computational Linguistics. http://aclweb.org/anthology/N12-1057.

Rao, Delip, David Yarowsky, Abhishek Shreevats, and Manaswi Gupta. 2010. "Classifying Latent User Attributes in Twitter." In *Proceedings of the 2nd International Workshop on Search and Mining User-Generated Contents,* edited by Jose Carlos Cortizo, Francisco M. Carrero, Iván Cantador, Jose Antonio Troyano, and Paolo Rosso, *26–30 October 2010, Toronto, Canada*, 37–44. New York: Association for Computing Machinery. doi:10.1145/1871985.1871993.

Rosenthal, Sara, and Kathleen McKeown. 2011. "Age Prediction in Blogs: A Study of Style, Content, and Online Behavior in Pre-and Post-Social Media Generations." In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, 1*9–24 June 2011, Portland, Oregon, USA,* program chairs Yuji Matsumoto and Rada Mihalcea, 763–72. Pennsylvania: Association for Computational Linguistics. http://anthology.aclweb.org/P/P11/P11-1077.pdf.

Schler, Jonathan, Moshe Koppel, Shlomo Argamon, and James W Pennebaker. 2006. "Effects of Age and Gender on Blogging." In *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*, 6:199–205. California: Association for the Advancement of Artificial Intelligence.

Silverstein, Michael. 2003. "Indexical Order and the Dialectics of Sociolinguistic Life." *Language & Communication* 23 (3): 193–229. doi:10.1016/S0271-5309(03)00013-2.

Tagliamonte, Sali. 2006. *Analysing Sociolinguistic Variation*. Cambridge: Cambridge University Press.

Tagliamonte, Sali A, and Derek Denis. 2008. "Linguistic Ruin? LOL! Instant Messaging and Teen Language." *American Speech* 83 (1): 3–34. doi: 10.1215/00031283-2008-001.

Thurlow, Crispin. 2006. "From Statistical Panic to Moral Panic: The Metadiscursive Construction and Popular Exaggeration of New Media Language in the Print Media." *Journal of Computer-Mediated Communication* 11 (3): 667–701. doi:10.1111/j.1083-6101.2006.00031.x.

Trudgill, Peter. 1972. "Sex, Covert Prestige and Linguistic Change in the Urban British English of Norwich." *Language in Society* 1 (02): 179–95. Cambridge: Cambridge University Press.

Walther, Joseph B, and Kyle P D'Addario. 2001. "The Impacts of Emoticons on Message Interpretation in Computer-Mediated Communication." In *Social Science Computer Review,* edited by David Garson, 19 (3): 324–47. doi: 10.1177/089443930101900307.

Werry, Christopher C. 1996. “Internet Relay Chat.” In *Computer-Mediated Communication: Linguistic, Social and Cross-Cultural Perspectives*, edited by Susan C. Herring, 47–63. Amsterdam/Pennsylvania: John Benjamins.

Wing, Benjamin P, and Jason Baldridge. 2011. “Simple Supervised Document Geolocation with Geodesic Grids.” In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, 1*9–24 June 2011*, *Portland, Oregon, USA*, program chairs Yuji Matsumoto and Rada Mihalcea, 955–64. Pennsylvania: Association for Computational Linguistics. http://www.aclweb.org/anthology/P11-1096.

Acknowledgments

We thank the reviewers of this submission for their helpful suggestions on improving the paper, and also the anonymous reviewers of the Joint Workshop on Social Dynamics and Personal Attributes in Social Media (at ACL 2014), who provided feedback at an early stage of this work. Thanks also to Brendan O'Connor for assistance with collecting the data, and to Adam Glynn for helpful discussions on causal inference. This research was supported by the National Science Foundation under award IIS-1111142.